

GENERALIZATIONS OF FOURIER ANALYSIS, AND HOW TO APPLY THEM

W.T. GOWERS

1. INTRODUCTION

This year's Colloquium Lectures are about an important theme in additive combinatorics. Additive combinatorics is a fairly recent branch of mathematics, though it grew out of much older ones. These notes are intended as an informal document to accompany the lectures and give more detail than it will be possible to give in the lectures themselves, though not full details about everything that is discussed. One aspect of its informality is that I have not provided a bibliography and many of the results I have mentioned, which are due to a variety of authors, are not attributed. In due course I will write a more formal version in which references will be provided. This will probably be published, and will definitely be posted on the arXiv.

Additive combinatorics is not very easy to characterize, but a good way to understand the flavour of the area is to look at one of its central theorems, the following famous result of Szemerédi from 1974, which solved a conjecture made by Erdős and Turán in 1936.

Theorem 1.1. *For every positive integer k and every $\delta > 0$ there exists a positive integer n such that every subset $A \subset \{1, 2, \dots, n\}$ of size at least δn contains an arithmetic progression of length k .*

This is a combinatorial theorem in the sense that we make no structural assumptions about A – it is just a subset of $\{1, 2, \dots, n\}$ of density at least δ . However, the set $\{1, 2, \dots, n\}$ has a rich additive structure, and that structure is highly relevant to the problem, since an arithmetic progression can be thought of as a sequence (x_1, x_2, \dots, x_k) such that

$$x_2 - x_1 = x_3 - x_2 = \dots = x_k - x_{k-1}.$$

(Of course, we also need to add the non-degeneracy condition that $x_1 \neq x_2$.)

However, there is more to additive combinatorics than a set of combinatorial theorems that involve addition in one way or another. To appreciate this, it is helpful to look at

the following statement, which turns out to be an equivalent reformulation of Szemerédi's theorem. The equivalence is a reasonably straightforward exercise to prove.

Theorem 1.2. *For every positive integer k and every $\delta > 0$ there exists a constant $c > 0$ such that for every positive integer n and every function $f : \mathbb{Z}_n \rightarrow [0, 1]$ that averages at least δ we have the inequality*

$$\mathbb{E}_{x,d} f(x)f(x+d) \dots f(x+(k-1)d) \geq c.$$

Here \mathbb{Z}_n is the cyclic group of order n and the notation $\mathbb{E}_{x,d}$ means the average over all x and d – that is, it is another way of writing $n^{-2} \sum_{x,d}$.

As n gets large, \mathbb{Z}_n is a better and better discrete approximation to the circle \mathbb{T} , which we can think of as the Abelian group consisting of all complex numbers of modulus 1. It is not hard to prove that the discrete statement is equivalent to the following continuous version.

Theorem 1.3. *For every positive integer k and every $\delta > 0$ there exists a constant $c > 0$ such that for every measurable function $f : \mathbb{T} \rightarrow [0, 1]$ that averages at least δ we have the inequality*

$$\mathbb{E}_{x,d} f(x)f(x+d) \dots f(x+(k-1)d) \geq c.$$

This time $\mathbb{E}_{x,d}$ stands for the integral with respect to the Haar measure on \mathbb{T}^2 .

This last reformulation illustrates an important point about many of the theorems of additive combinatorics (and extremal combinatorics more generally), which is that although they are combinatorial, they are also analytic. In fact, the more one thinks about them, the less important the distinction between discrete and continuous seems to be. And it is not just the statements that are (or can be made to be) analytic: a characteristic feature of much of additive combinatorics is that the proofs of its theorems use methods from areas of analysis such as functional analysis, Fourier analysis, and ergodic theory.

Here we shall focus on the last of these. Fourier analysis is an extremely useful tool for additive problems, and one of the aims of the Colloquium Lectures will be to explain why. Another aim, which is in some ways even more interesting, will be to demonstrate the *limitations* of Fourier analysis – that is, to look at problems that do not immediately yield to a Fourier-analytic approach. Sometimes that just means that one needs to look for a completely different kind of argument. However, with some problems the best way to make progress is not to abandon Fourier analysis altogether, but to generalize it in a suitable, and not always obvious, way. Thus, it sometimes happens that the limitations of one type of Fourier analysis lead to the development of another.

2. DISCRETE FOURIER ANALYSIS

Let $f : \mathbb{Z}_n \rightarrow \mathbb{C}$. We define its *discrete Fourier transform* $\hat{f} : \mathbb{Z}_n \rightarrow \mathbb{C}$ by the formula

$$\hat{f}(r) = \mathbb{E}_x f(x) \omega^{-rx},$$

where $\omega = \exp(2\pi i x/n)$ is a primitive n th root of unity. Note that there is a close resemblance between this formula, which we could equally well write as

$$\hat{f}(r) = \mathbb{E}_x f(x) \exp(-2\pi i r x/n),$$

and the familiar formulae for Fourier coefficients and Fourier transforms in the continuous setting. Of course, this is to be expected. Note also that the number ω^{-rx} is well-defined, since if r and n are integers, then adding a multiple of n to either of them makes no difference to it.

Although \hat{f} can be thought of as a function defined on \mathbb{Z}_n , it is more correct to regard it as defined on the dual group $\hat{\mathbb{Z}}_n$, which happens to be (non-naturally) isomorphic to \mathbb{Z}_n . The distinction has some importance in additive combinatorics, because the natural measures we put on \mathbb{Z}_n and $\hat{\mathbb{Z}}_n$ are different: for \mathbb{Z}_n we use the uniform probability measure, whereas for $\hat{\mathbb{Z}}_n$ we use the counting measure. This difference feeds into the definitions of some key concepts such as inner products, p -norms and convolutions. Given functions $f, g : \mathbb{Z}_n \rightarrow \mathbb{C}$ and $1 \leq p \leq \infty$, we have the following definitions.

- $\langle f, g \rangle = \mathbb{E}_x f(x) \overline{g(x)}$.
- $\|f\|_p = (\mathbb{E}_x |f(x)|^p)^{1/p}$.
- $f * g(x) = \mathbb{E}_{y+z=x} f(y)g(z)$.

The corresponding definitions for functions $\hat{f}, \hat{g} : \hat{\mathbb{Z}}_n \rightarrow \mathbb{C}$ are the same, but with sums replacing expectations. That is, they are as follows.

- $\langle \hat{f}, \hat{g} \rangle = \sum_x \hat{f}(x) \overline{\hat{g}(x)}$.
- $\|\hat{f}\|_p = (\sum_x |\hat{f}(x)|^p)^{1/p}$.
- $\hat{f} * \hat{g}(x) = \sum_{y+z=x} \hat{f}(y) \hat{g}(z)$.

With these measures in place, the familiar properties of the Fourier transform hold for the discrete Fourier transform as well, and have easier proofs. In particular, constant use is made of the following five rules, of which the first two are equivalent. All five are easy exercises.

- $\langle f, g \rangle = \langle \hat{f}, \hat{g} \rangle$ (Parseval's identity).
- $\|f\|_2 = \|\hat{f}\|_2$ (also Parseval's identity).

- $f(x) = \sum_r \hat{f}(r)\omega^{rx}$ (the inversion formula).
- $\widehat{f * g}(r) = \hat{f}(r)\hat{g}(r)$ (the convolution identity).
- If a is invertible mod n and $g(x) = f(ax)$ for every $x \in \mathbb{Z}_n$, then $\hat{g}(r) = \hat{f}(a^{-1}r)$ for every r (the dilation rule).

In additive combinatorics, one often deals with characteristic functions of subsets A of \mathbb{Z}_n , and some authors like to use the letter A for its own characteristic function: that is, $A(x) = 1$ if $x \in A$ and 0 otherwise. Given a subset $A \subset \mathbb{Z}_n$, define its *density* to be $|A|/n$. The following three observations are often used.

- $\hat{A}(0) = \alpha$.
- $\sum_r |\hat{A}(r)|^2 = \alpha$.
- $\hat{A}(-r) = \overline{\hat{A}(r)}$.

The first observation is immediate from the definition, the second follows from Parseval's identity and the fact that $\sum_r |\hat{A}(r)|^2 = \|\hat{A}\|_2^2$, and the third follows from the fact that A is real-valued and that $\omega^{rx} = \overline{\omega^{-rx}}$ for every r and x (and so is true of all real-valued functions).

3. ROTH'S THEOREM

To give an idea of how useful these simple facts are, we shall now sketch a proof of Roth's theorem, which is the case $k = 3$ of Szemerédi's theorem (Theorem 1.1 of these notes). Thus, we would like to prove the following theorem.

Theorem 3.1. *For every $\delta > 0$ there exists n such that every subset $A \subset \{1, 2, \dots, n\}$ of density at least δ contains an arithmetic progression of length 3.*

In order to apply Fourier analysis, it is convenient to think of A as a subset of \mathbb{Z}_n rather than of $\{1, 2, \dots, n\}$. (This is not essential, however: Roth originally treated A as a subset of \mathbb{Z} .) We shall also assume that n is odd. Let us write A_2 for the function defined by $A_2(z) = A(z/2)$, which is the characteristic function of the set of z such that $z/2 \in A$. Because n is odd, the map $z \mapsto z/2$ is a well-defined bijection.

The key observation that shows why Fourier analysis is useful is that the number of arithmetic progressions in A can be expressed in terms of convolutions, inner products and dilations, and therefore has a neat expression in terms of the Fourier coefficients of A .

Indeed, using the rules given earlier, we have that

$$\begin{aligned}
\mathbb{E}_{x+y=2z} A(x)A(y)A(z) &= \mathbb{E}_{x+y=z} A(x)A(y)A(z/2) \\
&= \mathbb{E}_z A * A(z)A_2(z) \\
&= \langle A * A, A_2 \rangle \\
&= \langle \widehat{A * A}, \hat{A}_2 \rangle \\
&= \langle \hat{A}^2, \hat{A}_2 \rangle \\
&= \sum_r \hat{A}(r)^2 \overline{\hat{A}_2(r)} \\
&= \sum_r \hat{A}(r)^2 \overline{\hat{A}(2r)} \\
&= \sum_r \hat{A}(r)^2 \hat{A}(-2r).
\end{aligned}$$

Why should this be useful? To answer that question, we need to bring in another simple but surprisingly powerful tool: the Cauchy-Schwarz inequality. First, recalling that $\hat{A}(0)$ is equal to the density of A , which we shall denote by α , we split the last expression up as

$$\alpha^3 + \sum_{r \neq 0} \hat{A}(r)^2 \hat{A}(-2r).$$

Thus, we have shown that

$$\mathbb{E}_{x+y=2z} A(x)A(y)A(z) = \alpha^3 + \sum_{r \neq 0} \hat{A}(r)^2 \hat{A}(-2r).$$

The left-hand side of this expression is the probability that x, y, z all belong to A if you choose them randomly to satisfy the equation $x + y = 2z$. Without the constraint that $x + y = 2z$ this probability would be α^3 , since each of x, y and z would have a probability α of belonging to A . So the term α^3 on the right-hand side can be thought of as “what one would expect” and the remainder of the right-hand side is a measure of the effect of the dependence of x, y and z on each other.

It is to bound the remainder term that we use the Cauchy-Schwarz inequality, and also the even more elementary inequality $|\langle f, g \rangle| \leq \|f\|_1 \|g\|_\infty$. We find that

$$\begin{aligned}
\left| \sum_{r \neq 0} \hat{A}(r)^2 \hat{A}(-2r) \right| &\leq \max_{r \neq 0} |\hat{A}(r)| \sum_{r \neq 0} |\hat{A}(r)| |\hat{A}(-2r)| \\
&\leq \max_{r \neq 0} |\hat{A}(r)| \left(\sum_r |\hat{A}(r)|^2 \right)^{1/2} \left(\sum_r |\hat{A}(-2r)|^2 \right)^{1/2} \\
&= \max_{r \neq 0} |\hat{A}(r)| \|\hat{A}\|_2^2 \\
&= \alpha \max_{r \neq 0} |\hat{A}(r)|.
\end{aligned}$$

It follows that

$$\mathbb{E}_{x+y=2z} A(x)A(y)A(z) \geq \alpha^3 - \alpha \max_{r \neq 0} |\hat{A}(r)|.$$

We see from this that if all the Fourier coefficients $\hat{A}(r)$ are small (more precisely, they have size significantly less than α^2), then the number of triples $(x, y, z) \in A^3$ with $x + y = 2z$ is close to $\alpha^3 n^2$, which is the approximate number we would get if the elements of A were chosen independently at random, each with probability α .

Therefore, either we have the arithmetic progression we are looking for (strictly speaking, this is incorrect because our triples satisfy the equation $x + y = 2z$ in \mathbb{Z}_n and not necessarily in \mathbb{Z} when we regard x, y and z as ordinary integers, but this is a technical problem that can be dealt with), or A has a large Fourier coefficient $\hat{A}(r)$ for some non-zero r . Here, “large” here can be taken to mean “of absolute value at least $c\alpha^2$ ” for some absolute constant $c > 0$.

In the second case, let us define a function $f : \mathbb{Z}_n \rightarrow \mathbb{R}$ by setting $f(x) = A(x) - \alpha$ for each x . It is easy to show that $\hat{f}(r) = \hat{A}(r)$ (this uses the fact that $r \neq 0$). So we obtain an inequality

$$|\hat{f}(r)| = |\mathbb{E}_x f(x) \omega^{-rx}| \geq c\alpha^2.$$

At this point we use a lemma, which I shall state imprecisely.

Lemma 3.2. *For every $r \neq 0$ there exists a partition of \mathbb{Z}_n into arithmetic progressions P_1, \dots, P_m , each of length at least $c\sqrt{n}$, such that the function ω^{rx} is approximately constant on each P_i .*

The proof of the lemma is an exercise based on a well-known technique: one uses the fact that by the pigeonhole principle it is possible to find $0 \leq u < v$ such that v is not too

large and $|\omega^{ru} - \omega^{rv}| = |1 - \omega^{r(v-u)}|$ is small. One can then partition \mathbb{Z}_n into arithmetic progressions of common difference $v - u$.

Given the lemma, one observes that

$$c\alpha^2 n \leq \left| \sum x f(x) \omega^{-rx} \right| \leq \sum_i \left| \sum_{x \in P_i} f(x) \omega^{-rx} \right| \approx \sum_i \left| \sum_{x \in P_i} f(x) \right|,$$

and also that

$$0 = \sum_x f(x) = \sum_i \sum_{x \in P_i} f(x).$$

Adding these equations together and using an averaging argument, we find that there exists i such that

$$\left| \sum_{x \in P_i} f(x) \right| + \sum_{x \in P_i} f(x) \geq c'\alpha^2 |P_i|,$$

where c' is a slightly smaller absolute constant (because of the approximation in the first equation), which implies that

$$\sum_{x \in P_i} f(x) \geq c'\alpha^2 |P_i|.$$

Recalling that $f(x) = A(x) - \alpha$ for each x , we find that this is telling us that

$$|A \cap P_i| \geq (\alpha + c'\alpha^2) |P_i|.$$

Thus, what we have managed to do is find an arithmetic progression P_i of length at least $c\sqrt{n}$ such that the density of A inside P_i is greater than the density of A inside \mathbb{Z}_n by $c'\alpha^2$.

We can iterate this argument: either $A \cap P_i$ contains an arithmetic progression of length 3 or P_i contains a subprogression of length at least $c\sqrt{|P_i|}$ inside which A has density at least $\alpha + 2c'\alpha^2$, and so on. The iteration must eventually terminate, because the density cannot exceed 1, and Roth's theorem is proved.

If one analyses carefully the bound that comes out of the above argument, one finds that it shows that if A is a subset of $\{1, 2, \dots, n\}$ of density at least $C/\log \log n$, for some absolute constant C , then A must contain an arithmetic progression of length 3. This bound has been improved in interesting ways several times. While these improvements are not the topic of these notes, it would be wrong not to mention them at all. The following table gives an idea of how the bounds have progressed over the years. The publication dates of the papers of Szemerédi and Heath Brown are slightly misleading: those results were actually independent. Also, the papers of Sanders obviously came out in the opposite order to the order in which the results were proved. The problem of improving the bounds for Roth's theorem has been an extremely fruitful one: the 2008 paper of Bourgain and

Bounds for Roth's theorem

Author	Density bound	Published
Roth	$C/\log \log n$	1953
Heath-Brown	$C/(\log n)^c$, some $c > 0$	1987
Szemerédi	$C/(\log n)^{1/20}$	1990
Bourgain	$C(\log \log n / \log n)^{1/2}$	1999
Bourgain	$C(\log \log n)^2 / (\log n)^{2/3}$	2008
Sanders	$(\log n)^{-3/4+o(1)}$	2012
Sanders	$C(\log \log n)^6 / \log n$	2011
Bloom	$C(\log \log n)^4 / \log n$	2012

the 2012 paper of Sanders could perhaps be regarded as clever refinements of existing techniques, but all the other papers introduced significant new ideas, many of which have been very influential and led to the solutions of several other problems.

To put these results in perspective, it is worth mentioning that the best known lower bound on the density (that is, the largest density known to be possible for a set that contains no progression of length 3) is $\exp(-c\sqrt{\log n})$, which is far lower than Bloom's current record upper bound. But even if that gap turns out to be very hard to close, we are tantalizingly close to a bound of $1/\log n$, which would be enough to give a purely combinatorial proof that the primes contain infinitely many arithmetic progressions of length 3 (a result that was proved by number-theoretic methods soon after Vinogradov proved his 3-primes theorem). In fact, a bound of $c \log \log n / \log n$ would suffice for this, since the fact that the primes have very small intersection with some arithmetic progressions (such as the even numbers) can be used to show that there are arithmetic progressions of length n inside which the primes have at least that density.

4. A FIRST GENERALIZATION – TO ARBITRARY FINITE ABELIAN GROUPS

Many of the proof techniques that give us results about subsets of \mathbb{Z}_n work just as well in an arbitrary Abelian group. This turns out to be a very useful observation, as there are some Abelian groups, in particular the groups \mathbb{F}_p^n for fixed p and large n , where the proofs are much cleaner. So sometimes to work out the proof of a result about \mathbb{Z}_n it is a good strategy to prove an analogue for a group such as \mathbb{F}_3^n first and then work out how to modify the argument so that it works in \mathbb{Z}_n .

Recall the inversion formula for the Fourier transform on \mathbb{Z}_n , which states that

$$f(x) = \sum_r \hat{f}(r) \omega^{rx}.$$

If we write ω_r for the function $x \mapsto \omega^{rx}$, then we can write the formula in the slightly more abstract form

$$f = \sum_r \hat{f}(r) \omega_r,$$

which is showing us how to write f as a linear combination of the functions ω_r .

What is special about the functions ω_r ? The property that singles them out is that they are the *characters* of \mathbb{Z}_n , that is, the homomorphisms from \mathbb{Z}_n to \mathbb{C} . It turns out to be straightforward to generalize Fourier analysis to all finite Abelian groups G by decomposing functions $f : G \rightarrow \mathbb{C}$ as linear combinations of characters.

For this to work, we would like the characters to form an orthonormal basis, which they do, by a well-known argument. To see the orthonormality, let χ be a non-trivial character, let $y \in G$ be such that $\chi(y) \neq 1$, and observe that

$$\mathbb{E}_x \chi(x) = \mathbb{E}_x \chi(xy) = \chi(y) \mathbb{E}_x \chi(x),$$

from which it follows that $\mathbb{E}_x \chi(x) = 0$. But then if χ_1 and χ_2 are distinct characters, we have that

$$\langle \chi_1, \chi_2 \rangle = \mathbb{E}_x \chi_1(x) \overline{\chi_2(x)} = \mathbb{E}_x \chi_1(x) \chi_2(x)^{-1},$$

which is zero, since $\chi_1 \chi_2^{-1}$ is a non-trivial character.

Less elementary is the fact that the characters span G . For this one needs the structure theorem for finite Abelian groups, which gives us that G is a product of cyclic groups. We know that each cyclic group has a complete basis of characters, and the products of those characters form a basis of characters for the whole group, which gives us a complete set.

Given that the characters form an orthonormal basis, we can expand a function f as a linear combination $\sum_\chi \langle f, \chi \rangle \chi$. The coefficients $\langle f, \chi \rangle$ are called the Fourier coefficients of f and denoted $\hat{f}(\chi)$. That is, we have the formula

$$\hat{f}(\chi) = \mathbb{E}_x f(x) \overline{\chi(x)}$$

for the Fourier transform, and the statement that $f = \sum_\chi \langle f, \chi \rangle \chi$ is giving us our inversion formula

$$f(x) = \sum_\chi \hat{f}(\chi) \chi(x).$$

The fact that we are writing $\hat{f}(\chi)$ represents a slight change of notation from the \mathbb{Z}_n case, where we wrote $\hat{f}(r)$ instead of $\hat{f}(\omega_r)$. This emphasizes the fact that properly speaking the Fourier transform is defined on the dual group \hat{G} rather than on G . It happens that these two groups are isomorphic, but the isomorphism is not natural in the category-theoretic sense.

When G is the group \mathbb{F}_3^n , the characters take the form $\omega_r : x \mapsto \omega^{r \cdot x}$, where now r and x are elements of \mathbb{F}_3^n , $\omega = \exp(2\pi i/3)$, and $r \cdot x$ is shorthand for $\sum_{i=1}^n r_i x_i$. It was observed by Meshulam that Roth's proof of Roth's theorem has an analogue for subsets of \mathbb{F}_3^n , and that the proof is in fact considerably simpler in that context because there is no longer any need for the lemma about partitioning into arithmetic progressions on which a character is roughly constant. The theorem is as follows.

Theorem 4.1. *There is a constant C such that for every positive integer n , every subset $A \subset \mathbb{F}_3^n$ of density at least C/n contains distinct elements x, y, z such that $x + y + z = 0$.*

Note that in \mathbb{F}_3^n the equation $x + y + z = 0$ is equivalent to the equation $x + y = 2z$, so the analogy with Roth's theorem is very close. As for the proof, one gets in exactly the same way that either A looks random enough that it must contain an arithmetic progression or there is a non-zero r such that $\hat{A}(r)$ (or $\hat{A}(\omega_r)$ if you prefer) has magnitude at least $c\alpha^2$, where α is the density of A . In the second case, it is easy to show that A has density at least $\alpha + c'\alpha^2$ in at least one of the three sets $\{x : r \cdot x = i\}$ (where $i = 0, 1$ or 2). Since these sets are just subspaces of \mathbb{F}_3^n of codimension 1, we are then already in a position to iterate.

This argument illustrates very well why it can be fruitful to look at more general Abelian groups. Because the group \mathbb{F}_3^n has a rich set of cosets of subgroups – namely all the affine subspaces – it is very convenient for iterative arguments. This somehow allows one to focus on the “real issues”. In more general Abelian groups, and in particular with the cyclic groups \mathbb{Z}_n , one has to make do with subsets that are “subgroup-like”. Doing so is possible, but it creates technical problems that can make arguments hard work to write down and even harder to read.

The strongest known lower bound for this problem is, once again, far lower. The best method we know of is to look for an example $B \subset \mathbb{F}_3^k$ for some small k and then to use it to create a class of examples $A = B^r \subset \mathbb{F}_3^{kr}$. If B has density c^k and $n = kr$, then B^r has density $c^{kr} = c^n$. But the following question is still wide open.

Question 4.2. *Let c_n be the greatest possible density of a subset $A \subset \mathbb{F}_3^n$ that contains no three distinct elements x, y, z such that $x + y + z = 0$. Does there exist $\theta < 1$ such that $c_n \leq \theta^n$ for every n ?*

Until fairly recently, the best known upper bound was given by the simple argument outlined above. But in 2011 Bateman and Katz improved the bound to one of the form $C/n^{1+\epsilon}$ for fixed constants C and $\epsilon > 0$. This was a remarkable achievement, given how long the bound had stood still, but the gap that remains is still huge.

5. THE U^2 NORM

In the proof of Roth's theorem, we had a useful measure of the quasirandomness of a function, namely the size of its largest Fourier coefficient – the smaller that size, the more quasirandom the function. However, this measure has the disadvantage that there isn't an obvious physical-space interpretation of $\|\hat{f}\|_\infty$ – that is, an expression in terms of the values of f that does not mention the Fourier transform. Instead, one often prefers to use the measure $\|\hat{f}\|_4$. In the contexts we care about, these two quantities are roughly equivalent, since we have the trivial inequalities

$$\|\hat{f}\|_\infty^4 \leq \|\hat{f}\|_4^4 \leq \|\hat{f}\|_\infty^2 \|\hat{f}\|_2^2,$$

and we usually deal with functions f such that $\|\hat{f}\|_2 = \|f\|_2 \leq 1$. This tells us that $\|\hat{f}\|_\infty$ is small if and only if $\|\hat{f}\|_4$ is small (though if we pass from one equivalent statement to the other and back again, we obtain a worse constant of smallness than the one we started with).

The reason that $\|\hat{f}\|_4$ is nice is that

$$\|\hat{f}\|_4^4 = \sum_r |\hat{f}(r)|^4 = \langle \hat{f}^2, \hat{f}^2 \rangle = \langle f * f, f * f \rangle = \mathbb{E}_{x+y=z+w} f(x)f(y)\overline{f(z)f(w)},$$

where in the above argument we used the definition of the ℓ_4 norm, the definition of the inner product on \hat{Z}_n , Parseval's identity and the convolution identity, and the definition of convolutions and inner products in \mathbb{Z}_n . (It is also possible to prove the identity above using a direct calculation, but it is nicer to use the basic properties of the Fourier transform.)

Quadruples (x, y, z, w) with $x + y = z + w$ are the same as quadruples of the form $(x, x + a + b, x + a, x + b)$, so the final expression above can be written in the form

$$\mathbb{E}_{x,a,b} f(x)\overline{f(x+a)f(x+b)}f(x+a+b).$$

Since this equals $\|\hat{f}\|_4^4$, we find that it is possible to define a norm $\|f\|_{U^2}$ by the formula

$$\|f\|_{U^2} = (\mathbb{E}_{x,a,b} f(x) \overline{f(x+a)} \overline{f(x+b)} f(x+a+b))^{1/4}.$$

This may seem pointless, since it is just renaming the norm $f \mapsto \|\hat{f}\|_4$, but we use a different name to emphasize that we are using a purely physical-space definition. The great advantage of doing this is that it gives us an alternative definition that is sometimes easier to generalize than the definition in terms of Fourier coefficients.

A useful fact about the U^2 norm is that it satisfies a kind of Cauchy-Schwarz inequality. Let us define a generalized inner product by the formula

$$[f_1, f_2, f_3, f_4] = \mathbb{E}_{x,a,b} f_1(x) \overline{f_2(x+a)} \overline{f_3(x+b)} f_4(x+a+b).$$

Then $\|f\|_{U^2}^4 = [f, f, f, f]$. The inequality states that

$$[f_1, f_2, f_3, f_4] \leq \|f_1\|_{U^2} \|f_2\|_{U^2} \|f_3\|_{U^2} \|f_4\|_{U^2}.$$

We quickly sketch a proof. We have that

$$\begin{aligned} [f_1, f_2, f_3, f_4] &= \mathbb{E}_{x,y,a} f_1(x) \overline{f_2(x+a)} \overline{f_3(y)} f_4(y+a) \\ &= \mathbb{E}_a (\mathbb{E}_x f_1(x) \overline{f_2(x+a)}) \overline{(\mathbb{E}_y f_3(y) \overline{f_4(y+a)})} \\ &\leq (\mathbb{E}_a |\mathbb{E}_x f_1(x) \overline{f_2(x+a)}|^2)^{1/2} (\mathbb{E}_a |\mathbb{E}_y f_3(y) \overline{f_4(y+a)}|^2)^{1/2} \end{aligned}$$

by the usual Cauchy-Schwarz inequality. But this last expression is easily seen to be

$$[f_1, f_2, f_1, f_2]^{1/2} [f_3, f_4, f_3, f_4]^{1/2}.$$

Furthermore, we have the symmetry $[f_1, f_2, f_3, f_4] = [f_1, f_3, f_2, f_4]$, so we can rewrite the last expression as

$$[f_1, f_1, f_2, f_2]^{1/2} [f_3, f_3, f_4, f_4]^{1/2}.$$

Applying the argument again we find that

$$[f_1, f_1, f_2, f_2] \leq [f_1, f_1, f_1, f_1]^{1/2} [f_2, f_2, f_2, f_2]^{1/2},$$

and similarly for f_3 and f_4 , and from this the result follows.

This inequality gives us a generalized Minkowski inequality in just the way that the normal Cauchy-Schwarz inequality gives the normal Minkowski inequality. Indeed,

$$\begin{aligned}
\|f_0 + f_1\|_{U^2}^4 &= [f_0 + f_1, f_0 + f_1, f_0 + f_1, f_0 + f_1] \\
&= \sum_{\epsilon \in \{0,1\}^4} [f_{\epsilon_1}, f_{\epsilon_2}, f_{\epsilon_3}, f_{\epsilon_4}] \\
&\leq \sum_{\epsilon \in \{0,1\}^4} \|f_{\epsilon_1}\|_{U^2} \|f_{\epsilon_2}\|_{U^2} \|f_{\epsilon_3}\|_{U^2} \|f_{\epsilon_4}\|_{U^2} \\
&= (\|f_0\|_{U^2} + \|f_1\|_{U^2})^4.
\end{aligned}$$

We thus have a proof, entirely in physical space, that the U^2 norm is a norm.

If $A \subset \mathbb{Z}_n$, then we can measure the quasirandomness of A as follows. Let α be the density of A and write $A(x) = \alpha + f(x)$. Then by the loose equivalence of the ℓ_∞ and ℓ_4 norms of the Fourier coefficients, we have that A is quasirandom in a useful sense if $\|f\|_{U^2}$ is small. One can check easily that $\|A\|_{U^2}^4 = \alpha^4 + \|f\|_{U^2}^4$, so this is saying that $\|A\|_{U^2}$ is approximately equal to α^4 . But $\|A\|_{U^2}^4$ has a nice interpretation. Recall that it equals

$$\mathbb{E}_{x+y=z+w} A(x)A(y)A(z)A(w),$$

which is the probability, if you choose a random quadruple (x, y, z, w) such that $x + y = z + w$, that all of x, y, z and w lie in A . This we call the *additive quadruple density* of A . Thus, a set of density α has additive quadruple density at least α^4 , with near equality if it is quasirandom in a useful sense.

An important final remark is that one can also prove entirely in physical space that if A is quasirandom in this sense, then its arithmetic-progression density is roughly α^3 . Indeed, writing $A(x) = \alpha + f(x)$ again, and noting that if we pick a random triple (x, y, z) with $x + y = 2z$, then any two of x, y and z will be independent and uniformly distributed (always assuming that n is odd), we have that

$$\begin{aligned}
\mathbb{E}_{x+y=2z} A(x)A(y)A(z) &= \mathbb{E}_{x+y=2z} (\alpha + f(x))(\alpha + f(y))(\alpha + f(z)) \\
&= \alpha^3 + \mathbb{E}_{x+y=2z} f(x)f(y)f(z) \\
&= \langle f * f, f \rangle,
\end{aligned}$$

where $f_2(z) = f(z/2)$ for each z . But by Cauchy-Schwarz and the fact that f takes values of modulus at most 1,

$$\begin{aligned} |\langle f * f, f_2 \rangle|^2 &\leq \|f * f\|_2^2 \|f_2\|_2^2 \\ &\leq \mathbb{E}_{x+y=z+w} f(x)f(y)\overline{f(z)f(w)} \\ &= \|f\|_{U^2}^4. \end{aligned}$$

Therefore, if $\|f\|_{U^2}$ is small, then $\mathbb{E}_{x+y=2z} A(x)A(y)A(z) \approx \alpha^3$.

In due course, we shall see how the arguments given above are more amenable to generalization than the Fourier-analytic proof we gave earlier.

We close this section by remarking that the definition of the U^2 norm and the basic observations we have made about it work just as well in an arbitrary finite Abelian group.

6. GENERALIZATION TO MATRICES

Given a function f we can define a linear map T_f that takes a function g to the convolution $f * g$. That is, we have

$$T_f(g)(x) = \mathbb{E}_u f(x-u)g(u).$$

If we define a matrix M_f by $M_f(x, u) = f(x-u)$, then this formula becomes

$$T_f(g) = \mathbb{E}_u M(f, u)g(u),$$

which is just the normal formula for multiplying a matrix by a vector, except that instead of summing over u we have taken the expectation. It will be convenient, for the purposes of this section, to adopt a non-standard definition of matrix multiplication by using this normalization. That is, we will say that if A and B are two matrices, then

$$(AB)(x, z) = \mathbb{E}_y A(x, y)B(y, z).$$

Since

$$T_f T_g h = T_f(g * h) = f * (g * h) = (f * g) * h,$$

we get that $M_f M_g = M_{f * g}$ with this normalization.

Notice that

$$T_f(\omega_r)(x) = f * \omega_r(x) = \mathbb{E}_u f(u)\omega_r^{r(x-u)} = \hat{f}(r)\omega_r(x).$$

Thus, ω_r is an eigenvector of T_f with eigenvalue $\hat{f}(r)$.

A more conceptual way of seeing this is to note that by the convolution identity, the convolution of f with the function $g = \sum_r \hat{g}(r)\omega_r$ is the function $\sum_r \hat{f}(r)\hat{g}(r)\omega_r$, so with respect to the basis $\omega_0, \dots, \omega_{n-1}$ all convolution maps T_f are multipliers (that is, given by diagonal matrices).

These observations allow us to translate some of the concepts we have defined so far into matrix language. The Fourier coefficients of a function f become the eigenvalues of the matrix M_f . However, that is just the beginning. Let us write $a \otimes b$ for the rank-1 matrix with

$$(a \otimes b)(u, v) = a(u)b(v).$$

Note that if $a, b, f : \mathbb{Z}_n \rightarrow \mathbb{C}$, then

$$(a \otimes \bar{b})(f)(x) = a(x)\mathbb{E}_y f(y)\overline{b(y)} = a(x)\langle f, b \rangle.$$

Thus, the diagonalization of f is telling us that

$$M_f = \sum_r \hat{f}(r)\omega_r \otimes \overline{\omega_r},$$

since if we apply either side to the function ω_s we obtain $\hat{f}(s)\omega_s$.

We are now in a position to write down Parseval's identity in matrix terms. First, note that

$$\mathbb{E}_{x,y} |M_f(x, y)|^2 = \mathbb{E}_{x,y} |f(x - y)|^2 = \mathbb{E}_x |f(x)|^2 = \|f\|_2^2.$$

Therefore, by Parseval's identity, we find that

$$\mathbb{E}_{x,y} |M_f(x, y)|^2 = \sum_r |\hat{f}(r)|^2.$$

The left-hand side is the L_2 norm of the matrix entries of M_f , which is often known as the (normalized) *Hilbert-Schmidt norm*. And the right-hand side, though it appears to be expressed in terms of f , can be thought of as the sum of squares of the eigenvalues of M_f .

This connection can be generalized to all matrices that have an orthonormal basis u_1, \dots, u_n of eigenvectors. In that case we can write $M = \sum_i \lambda_i u_i \otimes \overline{u_i}$ and we find

that

$$\begin{aligned}
\mathbb{E}_{x,y} |M(x,y)|^2 &= \mathbb{E}_{x,y} \sum_{i,j} \lambda_i \bar{\lambda}_j u_i(x) \overline{u_i(y) u_j(x)} u_j(y) \\
&= \sum_{i,j} \lambda_i \bar{\lambda}_j \mathbb{E}_{x,y} u_i(x) \overline{u_i(y) u_j(x)} u_j(y) \\
&= \sum_{i,j} \lambda_i \bar{\lambda}_j |\langle u_i, u_j \rangle|^2 \\
&= \sum_i |\lambda_i|^2.
\end{aligned}$$

More generally still, if M does not have an orthonormal basis of eigenvectors, it will still have a *singular value decomposition*, that is, a decomposition of the form $\sum_i \lambda_i u_i \otimes \bar{v}_i$ where $(u_i)_1^n$ and $(v_i)_1^n$ are both orthonormal bases and the λ_i are non-negative real numbers. (The non-negativity can be obtained by multiplying the v_i by suitable scalars of modulus 1.) The above argument carries over with very little change, and we find that $\|M\|_2^2$ (that is, the square of the normalized Hilbert-Schmidt norm) is equal to the sum of the squares of the singular values.

As we have already made clear, this fact specializes to Parseval's identity when the matrix is the matrix M_f of a convolution operator T_f .

More importantly, singular values of matrices play a rather similar role in graph theory to the role played by Fourier coefficients in additive combinatorics. To see this, let us first find an analogue for matrices of the U^2 norm. Given the correspondence so far, it should be equal to the ℓ_4 norm of the singular values, and its fourth power should have a nice interpretation in terms of the matrix values. This does indeed turn out to be the case. An argument similar to the one just given for the Hilbert-Schmidt norm, but slightly more complicated, shows that

$$\sum_i |\lambda_i|^4 = \mathbb{E}_{x,y,a,b} M(x,y) \overline{M(x+a,y) M(x,y+b)} M(x+a,y+b).$$

Now the fourth root of the left-hand side is a well-known matrix norm – the fourth-power trace class norm. From this one can deduce that the fourth root of the right-hand side is a norm, which we write as $\|M\|_{\square}$ and call the *box norm* (because we are summing over aligned rectangles). But as with the U^2 norm, one can prove this fact directly by first

defining a generalized inner product for two-variable functions

$$[f_1, f_2, f_3, f_4] = \mathbb{E}_{x,y,a,b} f_1(x, y) \overline{f_2(x+a, y) f_3(x, y+b) f_4(x+a, y+b)},$$

using the Cauchy-Schwarz inequality to prove that

$$[f_1, f_2, f_3, f_4] \leq \|f_1\|_{\square} \|f_2\|_{\square} \|f_3\|_{\square} \|f_4\|_{\square},$$

and finally deducing that $\|f + g\|_{\square}^4 \leq (\|f\|_{\square} + \|g\|_{\square})^4$ in more or less the same way as we did for the U^2 norm.

After this it will come as no surprise to learn that the box norm specializes to the U^2 norm when the matrix is a Toeplitz matrix (that is, the matrix of a convolution operator). Indeed, we have that

$$\begin{aligned} \|M_f\|_{\square}^4 &= \mathbb{E}_{x,y,a,b} f(x-y) \overline{f(x+a-y) f(x-y-b) f(x+a-y-b)} \\ &= \mathbb{E}_{x,a,b} f(x) \overline{f(x+a) f(x-b) f(x+a-b)} \\ &= \mathbb{E}_{x,a,b} f(x) \overline{f(x+a) f(x+b) f(x+a+b)} \\ &= \|f\|_{U^2}^4. \end{aligned}$$

Of course, we could also have deduced this less directly by using the relationship between eigenvalues, Fourier coefficients, and the two norms.

Now let us take a graph G and let M be its adjacency matrix. (That is, $M(x, y) = 1$ if there is an edge from x to y and 0 otherwise.) Then the analogy between subsets of \mathbb{Z}_n (or more general finite Abelian groups) and matrices strongly suggests that the box norm $\|\cdot\|_{\square}$ should be a useful measure of quasirandomness. That is indeed the case. If G has density δ , meaning that $\mathbb{E}_{x,y} M(x, y) = \delta$, then a straightforward argument using the Cauchy-Schwarz inequality shows that $\|M\|_{\square} \geq \delta$. If equality almost holds, then G turns out to enjoy a number of properties that typical random graphs have.

To see this, we begin by noting that the box norm relates to the largest singular value in much the way that the U^2 norm relates to the largest Fourier coefficient. If the singular values are $\lambda_1, \dots, \lambda_n$ and if $\lambda = (\lambda_1, \dots, \lambda_n)$, then

$$\|\lambda\|_{\infty}^4 \leq \|\lambda\|_4^4 \leq \|\lambda\|_2^2 \|\lambda\|_{\infty}^2,$$

and if the matrix entries have modulus at most 1 then we know in addition that $\|\lambda\|_2^2 = \|M\|_2^2 \leq 1$. Therefore, the largest singular value (which is equal to the operator norm of the matrix) is small if and only if the box norm is small.

For convenience let us now assume that G is regular, so every vertex has degree δn . (The proofs become slightly more complicated if we do not have this.) Then the constant function $u(x) = 1$ is an eigenvector of M with eigenvalue δ . (Recall that we are using expectations in our matrix multiplication, which is why we get δ here rather than δn .)

Now consider the matrix $A = M - \delta u \otimes u$. That is, $A(x, y) = M(x, y) - \delta$. Since G is regular, we find that $\mathbb{E}_x A(x, y) = 0$ for every y and $\mathbb{E}_y A(x, y) = 0$ for every x . From this it is not hard to prove that $\|M\|_{\square}^4 = \delta^4 + \|A\|_{\square}^4$: we expand $\|A + \delta u \otimes u\|_{\square}^4$ as a sum of sixteen terms and the only ones that are not zero are the term with all A s and the term with all δ s.

Therefore, if $\|M\|_{\square}$ is close to δ , it follows that $\|A\|_{\square}$ is close to zero, which implies that the largest singular value of A is small, and therefore that A has a small operator norm. Let θ be this operator norm.

Now let f and g be two functions defined on the vertex set of G that take values in the interval $[-1, 1]$. Then

$$|\langle Af, g \rangle| \leq \|Af\|_2 \|g\|_2 \leq \theta \|f\|_2 \|g\|_2 \leq \theta.$$

We also have that

$$\langle (\delta u \otimes u)(f), g \rangle = \langle (\delta \mathbb{E}_x f(x))u, g \rangle = \delta \mathbb{E}_x f(x) \mathbb{E}_y g(y).$$

It follows that

$$|\langle Mf, g \rangle - \delta \mathbb{E}_x f(x) \mathbb{E}_y g(y)| \leq \theta.$$

But $\langle Mf, g \rangle = \mathbb{E}_{x,y} M(x, y) f(x) g(y)$, so if θ is small then this is telling us that

$$\mathbb{E}_{x,y} M(x, y) f(x) g(y) \approx \delta \mathbb{E}_{x,y} f(x) g(y).$$

Suppose now that f and g are the characteristic functions of sets U and V of density α and β . Now we have that

$$\mathbb{E}_{x,y} M(x, y) U(x) V(y) \approx \delta \mathbb{E}_{x,y} U(x) V(y) = \delta \alpha \beta.$$

This tells us that the number of edges from U to V in the graph is approximately $\delta |U| |V|$, which is exactly the number one would expect if G was a random graph with density δ .

Now the fourth power of the box norm of M can be seen to equal the *4-cycle density* of the graph G , that is, the probability, if vertices x_1, x_2, x_3, x_4 are chosen independently at random, that $x_1 x_2, x_2 x_3, x_3 x_4$ and $x_4 x_1$ are all edges of G . Thus, we have started with a ‘‘local’’ assumption – that the number of 4-cycles in the graph is almost as small as it can possibly be given the density of the graph – and ended up with a global conclusion – that

the number of edges between any two large sets is approximately what one would expect in a random graph of the same density. This fact has many applications in graph theory.

The converse can also be shown without too much difficulty. In fact, there turn out to be several properties that are all loosely equivalent and all say that in one way or another a graph G behaves like a random graph. A particularly interesting one from the point of view of comparison with Roth's theorem is the statement that if a graph G of density δ is quasirandom (in, for example, the sense of having box norm approximately δ) then for any graph H with k edges (here k is fixed and the size of G is tending to infinity) the H density in G is approximately δ^k , as it would be in a random graph. Conversely, if G contains the "wrong" number of copies of H , then we can find a subgraph that is substantially denser than the original graph.

It is worth pointing out that not all the basic properties of the Fourier transform carry over in a nice way to matrices. For example, the inner product corresponding to the normalized Hilbert-Schmidt norm is

$$\langle A, B \rangle = \mathbb{E}_{xy} A(x, y) \overline{B(x, y)} = \text{tr}(AB^*)$$

(where I have also defined the trace in a normalized way – that is, $\text{tr}(A) = \mathbb{E}_x A_{xx}$). If the singular-value decompositions of A and B are $\sum_i \lambda_i u_i \otimes \overline{v_i}$ and $\sum_j \mu_j w_j \otimes \overline{z_j}$, then $\langle A, B \rangle$ works out to be

$$\sum_{i,j} \lambda_i \overline{\mu_j} \langle u_i, w_j \rangle \overline{\langle v_i, z_j \rangle}.$$

If it happens that $u_i = w_i$ and $v_i = z_i$ for every i , as it does when $A = B$, then this simplifies to $\sum_i \lambda_i \overline{\mu_i}$, the formula we would like, but if not then we have to make do with the more complicated formula above (which nevertheless can be useful sometimes).

Similarly, there is no tidy analogue of the convolution identity except under very special circumstances. In general,

$$\left(\sum_i \lambda_i u_i \otimes \overline{v_i} \right) \left(\sum_j \mu_j w_j \otimes \overline{z_j} \right) = \sum_{i,j} \lambda_i \mu_j \overline{\langle w_j, v_i \rangle} u_i \otimes \overline{z_j}.$$

If $v_i = w_i$ for every i , then this simplifies to $\sum_i \lambda_i \mu_i u_i \otimes \overline{z_i}$, so we find that the singular values of the matrix product are products of the singular values of the original matrices. But this is an unusual situation (that happens to occur when the two matrices are convolution matrices and all the bases are the same basis of trigonometric functions).

7. QUADRATIC FOURIER ANALYSIS

In this section I shall discuss a generalization of Fourier analysis that lacks a satisfactory inversion formula. This might seem to be such a fundamental property of the Fourier transform that the generalization does not deserve to be called a generalization of Fourier analysis. However, for several applications of Fourier analysis, a weaker property suffices, and that weaker property can be generalized. It is, however, a very interesting open problem to develop the theory further so as to make the analogy with conventional discrete Fourier analysis closer. But first, let us look at a problem that demonstrates the need for a generalization at all, namely Szemerédi's theorem for progressions of length 4.

At the heart of the proof for progressions of length 3 is the identity

$$\mathbb{E}_{x+y=2z} f(x)g(y)h(z) = \sum_r \hat{f}(r)\hat{g}(r)\hat{h}(-2r).$$

We have essentially proved this already, but a variant of the argument is to observe that both sides are equal to $\mathbb{E}_{x,y,z} f(x)g(y)h(z) \sum_r \omega^{-r(x+y-2z)}$. So it is natural to look for a similar identity for progressions of length 4. Such a progression can be thought of as a quadruple (x, y, z, w) such that $x + z = 2y$ and $y + w = 2z$. However,

$$\begin{aligned} \mathbb{E}_{x+z=2y, y+w=2z} f_1(x)f_2(y)f_3(z)f_4(w) \\ &= \mathbb{E}_{x,y,z,w} f_1(x)f_2(y)f_3(z)f_4(w) \sum_{r,s} \omega^{-r(x-2y+z)-s(y-2z+w)} \\ &= \sum_{r,s} \hat{f}_1(r)\hat{f}_2(-2r+s)\hat{f}_3(r-2s)\hat{f}_4(s). \end{aligned}$$

A quadruple (a, b, c, d) can be written in the form $(r, -2r + s, r - 2s, s)$ if and only if $3a + 2b + c = b + 2c + 3d = 0$. So we have ended up with a sum over four variables that satisfy two linear equations, which is what we had before we took the Fourier transform. So we have not gained anything.

An even more compelling argument that the Fourier transform is too blunt a tool for our purposes is to note that it is possible for all the Fourier coefficients of f_1, f_2, f_3 and f_4 to be tiny, but for the expectation $\mathbb{E}_{x,d} f_1(x)f_2(x+d)f_3(x+2d)f_4(x+3d)$ to be large. (This is another way of writing the expression that we have just evaluated in terms of Fourier transforms.) Let $f_1(x) = f_4(x) = \omega^{x^2}$ and let $f_2(x) = f_3(x) = \omega^{3x^2}$. Then

$$\mathbb{E}_{x,d} f_1(x)f_2(x+d)f_3(x+2d)f_4(x+3d) = \mathbb{E}_{x,d} \omega^{x^2-3(x+d)^2+3(x+2d)^2-(x+3d)^2}.$$

But the exponent on the right-hand side is identically zero, so both sides are equal to 1, which is as large as the expectation can possibly be given that all four functions take values of modulus 1. On the other hand, functions like ω^{x^2} have tiny Fourier coefficients. To see this (assuming for convenience that n is odd), note that if $f(x) = \omega^{x^2}$, then

$$\hat{f}(r) = \mathbb{E}_x \omega^{x^2 - rx} = \mathbb{E}_x \omega^{(x-r/2)^2 - r^2/4} = \omega^{-r^2/4} \mathbb{E}_x \omega^{x^2}.$$

This shows that $|\hat{f}(r)| = |\mathbb{E}_x \omega^{x^2}|$ is the same for all r , and therefore by Parseval it equals $n^{-1/2}$ for all r . In other words, the largest Fourier coefficient is as small as Parseval's identity will allow.

It is almost impossible at this stage not to have the following thought. For Roth's theorem, the functions that caused trouble by not being sufficiently random-like were the trigonometric functions $x \mapsto \omega^{rx}$. These are *linear phase functions* – that is, compositions of linear functions with the function $x \mapsto \omega^x$. We have just seen that when it comes to discussing arithmetic progressions of length 4, quadratic phase functions, that is, functions of the form $\omega^{q(x)}$ where q is a quadratic, cause problems. Could it be that these are somehow the only functions that cause problems? Does there exist some kind of “quadratic Fourier analysis” that allows one to expand a function as a linear combination of quadratic phase functions and thereby to generalize the proof of Roth's theorem to progressions of length 4?

The answer to this question turns out to be a partial yes. More precisely, one can generalize “linear” Fourier analysis by just enough to obtain a proof of Szemerédi's theorem for progressions of length 4, but the generalized Fourier analysis lacks some of the nice properties of the usual Fourier transform, as a result of which the proof becomes substantially harder. In particular, it turns out that the quadratic phase functions are *not* the only ones that cause trouble – there are also some more general functions that exhibit sufficiently quadratic-like behaviour to cause problems similar to the ones caused by the “pure” quadratic phase functions. But before we get on to that, it will be useful to look at another concept that comes into the picture.

8. THE U^3 NORM

Discrete Fourier analysis decomposes a function into characters. It is far from obvious how to define a “quadratic” analogue of this decomposition, since one's natural first guesses turn out not to have the properties one wants, as we shall see later. But already it is clear that there are problems, because there are n^2 functions of the form $x \mapsto \omega^{ax^2+bx}$, so we cannot hope to define a quadratic Fourier transform by simply writing down a suitable basis of \mathbb{C}^n and expanding functions in terms of that basis.

It is for this reason that the reformulation of the norm $f \mapsto \|f\|_4$ in purely physical-space terms is so important. It gives us a concept that *is* easy to generalize. As one might expect, there are U^k norms for all $k \geq 2$ (and also a seminorm when $k = 1$), but it is clear what they are once one has seen the U^3 norm. It is defined by the formula

$$\|f\|_{U^3}^8 = \mathbb{E}_{x,a,b,c} f(x) \overline{f(x+a)} f(x+b) \overline{f(x+a+b)} f(x+c) \overline{f(x+a+c)} f(x+b+c) \overline{f(x+a+b+c)}$$

That is, where the U^2 norm involves an average over “squares”, the U^3 norm involves a similar average over “cubes”, and the U^k norm involves a similar average over k -dimensional cubes. The letter U stands for “uniformity”, because when a function has a small uniformity norm, its values are “uniformly distributed” in a useful sense.

There are a few remarks to make about the U^3 norm to give an idea of its basic properties and of why it is likely to be important to us.

- First, it really is a norm. This is proved in much the same way as it is for the U^2 norm: one defines an appropriate generalized inner product (by using eight different functions in the formula above instead of just one), deduces a generalized Cauchy-Schwarz inequality from the conventional Cauchy-Schwarz inequality, and finally a generalized Minkowski inequality from the generalized Cauchy-Schwarz inequality.
- Secondly, if f is a quadratic phase function $f(x) = \omega^{rx^2+sx}$, then $\|f\|_{U^3}$ takes the largest possible value (given that all the values of f have modulus 1), namely 1. This is simple to check, and boils down to the fact that

$$\begin{aligned} x^2 - (x+a)^2 - (x+b)^2 + (x+a+b)^2 - (x+c)^2 \\ + (x+a+c)^2 + (x+b+c)^2 - (x+a+b+c)^2 = 0 \end{aligned}$$

for every x, a, b and c .

- Thirdly, the U^k norms increase as k increases. In particular, the U^3 norm is larger than the U^2 norm. This means that the statement that $\|f\|_{U^3}$ is small is stronger than the statement that $\|f\|_{U^2}$ is small. That fact, combined with the observation that $\|f\|_{U^3}$ is large for quadratic phase functions, gives some reason to hope that the U^3 norm could be a useful measure of quasirandomness for Szemerédi’s theorem for progressions of length 4.

- Fourthly, if A is a set of density α , then an easy Cauchy-Schwarz argument shows that $\|A\|_{U^3} \geq \alpha$. Also, $\|A\|_{U^3}^8$ counts the number of “cubes” in A . So when we talk about sets, we will want to regard a set as “quadratically uniform” if it has almost the minimum number of cubes, and this will be a stronger property than the “linear uniformity” that we used in the proof of Roth’s theorem.

Presenting those remarks is slightly misleading, however, as it suggests that the definition of the U^3 norm is a purely speculative generalization of the definition of the U^2 norm that just happens to be useful. In fact, the definition arises naturally (or at least can arise naturally) when one tries to generalize the physical-space argument we saw earlier that shows that a set with small U^2 norm has roughly the expected number of arithmetic progressions of length 3. One ends up being able to show that if f_1, f_2, f_3 and f_4 are functions that take values of modulus at most 1, then

$$|\mathbb{E}_{x,d} f_1(x) f_2(x+d) f_3(x+2d) f_4(x+3d)| \leq \min_i \|f_i\|_{U^3}.$$

In other words, if one of the four functions has a small U^3 norm, then the arithmetic progression count must be small.

The point I am making here is that if one sets out to prove a bound for the left-hand side in terms of some suitable function of f_4 , say, knowing that one’s main tool is the Cauchy-Schwarz inequality, then the function that one obtains is precisely the U^3 norm.

The inequality above can be used to show that if A is a set of density α and $\|A\|_{U^3} \leq \alpha + c(\alpha)$, then A is sufficiently quasirandom to contain an arithmetic progression of length 4, and in fact to have 4-AP density approximately α^4 . To prove this, one writes $A = \alpha + f$ with $\|f\|_{U^3}$ small, one expands out the expression

$$\mathbb{E}_{x,d} A(x) A(x+d) A(x+2d) A(x+3d)$$

as a sum of 16 terms, and one uses the inequality above to show that all these terms are small apart from the main term α^4 .

9. GENERALIZED QUADRATIC PHASE FUNCTIONS

In the previous section we noted that if q is a quadratic function defined on \mathbb{Z}_n , and f is the function $f(x) = \omega^{q(x)}$, then $\|f\|_{U^3} = 1$, which is as large as it can possibly be. The key

to this fact, as we have already noted, is that quadratic functions have the property that

$$\begin{aligned} q(x) - q(x+a) - q(x+b) + q(x+a+b) - q(x+c) \\ + q(x+a+c) + q(x+b+c) - q(x+a+b+c) = 0 \end{aligned}$$

for every x, a, b, c . Moreover, this property characterizes quadratic functions.

However, if we do not insist on maximizing $\|f\|_{U^3}$ but merely getting *close* to the maximum, then we suddenly let in a whole lot more functions. In this section I shall describe one or two of them.

There is a general recipe for producing them, which is to take a set $A \subset \mathbb{Z}_n$ and construct a *quadratic homomorphism* on A – that is, a map $\psi : A \rightarrow \mathbb{C}$ that takes values of modulus 1 and satisfies the equation

$$\psi(x)\overline{\psi(x+a)}\overline{\psi(x+b)}\overline{\psi(x+a+b)}\overline{\psi(x+c)}\overline{\psi(x+a+c)}\overline{\psi(x+b+c)}\overline{\psi(x+a+b+c)} = 1$$

whenever all of $x, x+a, x+b, x+c, x+a+b, x+a+c, x+b+c$ and $x+a+b+c$ all belong to A . (As we have already noted, if A has density α , there will be at least $\alpha^8 n^4$ “cubes” of this kind.) We then define $f(x)$ to be $\psi(x)$ for $x \in A$ and 0 otherwise. For this to produce interesting examples, we need to choose our set A carefully, but that can be done.

As a first example, take A to be the set $\{1, \dots, \lfloor n/2 \rfloor\}$. If we now let β be any real number, we can define $f(x)$ to be $e^{2\pi i \beta x^2}$ on A and zero outside. If β is a multiple of $1/n$, then this will give us a function ω^{rx^2} restricted to A . However, if we choose β not to be close to a multiple of $1/n$ we can obtain functions that do not even correlate with functions of the form ω^{rx^2+sx} . Suppose, for example, that we take $\beta = 1/2n$. Then our best chance of a correlation will be with either the constant function 1 or the function $\omega^{x^2} = e^{4\pi i \beta x^2}$. In both cases, the inner product has modulus $n^{-1} |\sum_{x \in A} e^{\pi i x^2/n}|$, which can be shown to be small by a simple trick known as Weyl differencing: we observe that

$$\left| \sum_{x \in A} e^{\pi i x^2/n} \right|^2 = \sum_{x,y} e^{\pi i (x^2 - y^2)/n} = \sum_{x,y} e^{\pi i (x+y)(x-y)}.$$

The last sum can be split into a sum of geometric progressions, each of which can be evaluated explicitly, and almost all of which turn out to be small. Essentially the same technique proves that in fact our function f has a very small correlation with any function of the form $\omega^{q(x)}$ for a quadratic function q defined on \mathbb{Z}_n .

It is worth stopping to think about why a similar argument does not show that we have to consider more functions even in the linear case. What if we take a function on the set

A above of the form $e^{2\pi i\beta x}$ with β far from a multiple of $1/n$? In fact, what if we take $\beta = 1/2n$ as before?

In this case the correlation with a constant function has magnitude $n^{-1}|\sum_{x\in A} e^{\pi ix/n}|$, and x/n lies between 0 and $1/2$. It follows that all the numbers $e^{\pi ix/n}$ are on one side of the unit circle, and the result is that we do not get the cancellation that occurred with the quadratic example above. The difference between the two situations is that the function $e^{\pi ix^2/n}$ jumps round the circle many times, whereas the function $e^{\pi ix/n}$ does not – which is due to the fact that the function x^2 grows much more rapidly than the function x .

Another way of choosing a set A is to make it look like a portion of \mathbb{Z}^d for some small d . To give an example with $d = 2$, let $m = \lfloor \sqrt{n}/2 \rfloor$ and let A consist of all numbers of the form $x + 2my$ such that $x, y \in \{0, 1, \dots, m-1\}$. This we can think of as a two-dimensional set with basis 1 and $2m$: the pair (x, y) then represents the point $x + 2my$ in coordinate form.

An obvious class of functions to take on a multidimensional set is the class of quadratic forms, and we can do that here. We pick coefficients $a, b, c \in \mathbb{Z}_n$ and define $f(x + 2my)$ to be $\omega^{ax^2 + bxy + cy^2}$ for all $x, y \in \{0, 1, \dots, m-1\}$ and take all other values of f to be zero. It is easy to check that f is a quadratic homomorphism in the sense just defined, and it can also be shown that f does not correlate with any pure quadratic phase function.

We can of course combine these ideas by taking more general coefficients. We can also define a wide variety of two-dimensional sets by taking different “basis vectors”, and we can increase the dimension. Thus, the set of functions we are forced to consider is much richer than the corresponding set for the U^2 norm.

10. SZEMERÉDI’S THEOREM FOR PROGRESSIONS OF LENGTH 4

We remarked at the end of Section 8 that if a set A is quasirandom in the sense of having an almost minimal U^3 norm, then it contains an arithmetic progression of length 4. Furthermore, the proof of this fact is closely analogous to the proof of the corresponding fact relating the U^2 norm to arithmetic progressions of length 3. So it is natural to try to continue the analogy and complete a proof of Szemerédi’s theorem for progressions of length 4. That is, we would like to argue that if the U^3 norm of A is *not* approximately minimal, then we can obtain a density increase on an appropriate subspace.

At this point we find that we are a little stuck. In the U^2 case we used the fact that if f is a function taking values of modulus at most 1, and $\|f\|_{U^2} = \|\hat{f}\|_4$ is bounded below by a positive constant c , then $\|\hat{f}\|_\infty$ is bounded below by c^2 , which we can use to argue

that a set with no arithmetic progression of length 3 must be sufficiently “unrandom” to correlate well with a trigonometric function. So to continue the analogy, it looks as though we need to find norms $\|\cdot\|$ and $|||\cdot|||$ with the following properties.

- (1) The norm $\|\cdot\|$ is defined in a different way from the U^3 norm, but happens to be equal to it.
- (2) If $\|f\|_\infty \leq 1$ and $\|f\| \geq c$, then one can prove very straightforwardly that $|||f||| \geq \gamma(c)$ (where $\gamma(c) > 0$ if $c > 0$, and ideally the dependence will be a good one).
- (3) The fact that $|||f||| \geq \gamma$ is telling us that there is some function $\psi \in \Psi$ for which $|\langle f, \psi \rangle| \geq \theta(\gamma)$, where Ψ is a class of “nice” functions (which will probably exhibit behaviour similar to that of quadratic phase functions).
- (4) If A is a set of density α , $f = A - \alpha$, and $|\langle f, \psi \rangle| \geq \theta$ for some $\psi \in \Psi$, then there is a long subprogression P inside which A has density at least $\alpha + \eta(\theta)$.

Implicit in the third of these conditions is that $|||\cdot|||$ and Ψ are related by the formula $\max\{|\langle f, \psi \rangle| : \psi \in \Psi\}$.

The big problem we face is that there is no obvious reformulation of the U^3 norm analogous to the reformulation $\|f\|_{U^2} = \|\hat{f}\|_4$ of the U^2 norm. So we do not know of a candidate for $\|\cdot\|$. However, that does not mean that there is nothing we can do, since there is still the possibility of passing directly from the statement that $\|f\|_{U^3} \geq c$ to the statement that $|\langle f, \psi \rangle| \geq \theta(c)$ for some suitably nice function ψ , or even bypassing this statement and heading straight for the conclusion that A is denser in some long subprogression. Both approaches turn out to be possible.

It is not possible here to do more than give a very brief sketch of how the proof works. We start with a function f with $\|f\|_\infty \leq 1$ and $\|f\|_{U^3}^8 \geq \gamma$. That inequality expands to the inequality

$$\mathbb{E}_{x,a,b,c} f(x) \overline{f(x-a)} \overline{f(x-b)} \overline{f(x-a-b)} \overline{f(x-c)} \\ f(x-a-c) f(x-b-c) \overline{f(x-a-b-c)} \geq \gamma,$$

where we have switched from plus signs to minus signs for unimportant aesthetic reasons. We now define, for each a , a function $\partial_a f$ by the formula $\partial_a f(x) = f(x) \overline{f(x-a)}$, which allows us to rewrite the inequality above as

$$\mathbb{E}_a \mathbb{E}_{x,b,c} \partial_a f(x) \overline{\partial_a f(x-b)} \overline{\partial_a f(x-c)} \partial_a f(x-b-c) \geq \gamma.$$

Now this is just telling us that $\mathbb{E}_a \|\partial_a f\|_{U^2}^4 \geq \theta$, from which it follows that there must be several a for which $\|\partial_a f\|_{U^2}$ is large. By the rough equivalence of the U^2 norm with the magnitude of the largest Fourier coefficient, we can deduce from this that several of the functions $\partial_a f$ have at least one large Fourier coefficient. It follows that there is a large set B and a function $\phi : B \rightarrow \mathbb{Z}_n$ such that $\widehat{\partial_a f}(\phi(a))$ is large for every $a \in B$. More formally, we can obtain an inequality

$$\mathbb{E}_a B(a) |\widehat{\partial_a f}(\phi(a))|^2 \geq \theta$$

for some θ that depends (polynomially) on γ .

It turns out that one can perform some algebraic manipulations with this statement and eventually prove that the function ϕ has an interesting “partial additivity” property, which states that there are at least ηn^3 quadruples $(x, y, z, w) \in B^4$ (for some η that depends on γ only) such that

$$x + y = z + w$$

and

$$\phi(x) + \phi(y) = \phi(z) + \phi(w).$$

This property appears at first to be somewhat weak, since it tells us that ϕ is additive on only a small percentage of the quadruples $x + y = z + w$. Remarkably, however, this is another instance where a local assumption can be used to prove a global conclusion: the only way that ϕ can be this additive is if it has a form that can be described very precisely.

Recall the two-dimensional set we defined in the previous section. It is an example of a *two-dimensional arithmetic progression*. More generally, a k -dimensional arithmetic progression is a set of the form

$$\{x + a_1 d_1 + a_2 d_2 + \cdots + a_k d_k : 0 \leq a_i < m_i\}.$$

The numbers d_1, \dots, d_k are the common differences and the numbers m_1, \dots, m_k are the lengths. The arithmetic progression is called *proper* if it has cardinality $m_1 \dots m_k$ – that is, no two of the $a_1 d_1 + \cdots + a_k d_k$ coincide.

Given such a progression and coefficients $\mu_0, \mu_1, \dots, \mu_k \in \mathbb{Z}_n$ one can define something like a linear form by the obvious formula

$$x + a_1 d_1 + a_2 d_2 + \cdots + a_k d_k \mapsto \mu_0 + \sum_i \mu_i a_i.$$

Let us call such a map *quasilinear*.

The result that tells us about the structure of ϕ is the following.

Theorem 10.1. *For every $\eta > 0$ there is an integer $d = d(\eta)$ and a constant $\zeta = \zeta(\eta) > 0$ with the following properties. Let $B \subset \mathbb{Z}_n$ and suppose that there are ηn^3 quadruples $(x, y, z, w) \in B^4$ with $x + y = z + w$ and $\phi(x) + \phi(y) = \phi(z) + \phi(w)$. Then there is a proper arithmetic progression P of dimension at most d and a quasilinear map $\psi : P \rightarrow \mathbb{Z}_n$ such that for at least ζn values of $x \in \mathbb{Z}_n$ we have that $x \in B \cap P$ and $\phi(x) = \psi(x)$.*

Loosely speaking, this tells us that there must be a quasilinear map that agrees a lot of the time with ϕ . To prove this, one must use some important results in additive combinatorics, such as a famous theorem of Freiman (and more particularly a proof of the theorem due to Ruzsa) as well as a quantitative version of a theorem of Balog and Szemerédi.

Now let us see why it is plausible that linear behaviour of the function ϕ should lead to quadratic behaviour in the function f from which it was derived. Consider an example where f is defined by a formula of the form $f(x) = \omega^{\nu(x)}$. Then $\partial_a f(x) = \omega^{\nu(x) - \nu(x-a)}$. So the statement that $\widehat{\partial_a f}(\phi(a))$ is large is telling us that the functions $\omega^{\nu(x) - \nu(x-a)}$ and $\omega^{a\phi(x)}$ correlate well. Since ϕ exhibits linear behaviour, the function $(a, x) \mapsto a\phi(x)$ exhibits bilinear behaviour.

But that is exactly what happens when ν is a quadratic function: if $\nu(x) = rx^2 + sx$, then $\nu(x) - \nu(x-a) = 2rxa - ra^2 + sa$, which implies that $\partial_a f$ has a large Fourier coefficient at $2ra$.

At this point one can use the information we have in a reasonably straightforward way to prove a weakish statement that is sufficient for Szemerédi's theorem, or we can work harder to prove a stronger statement that can be thought of as giving us some kind of quadratic Fourier analysis. The weakish statement (stated qualitatively) is the following.

Lemma 10.2. *Let $f : \mathbb{Z}_n \rightarrow \mathbb{C}$ be a function with $\|f\|_\infty \leq 1$ and suppose that there exists a quasilinear function ψ defined on a low-dimensional arithmetic progression P such that $\widehat{\partial_a f}(\psi(a))$ is large for many $a \in P$. Then there are long arithmetic progressions P_1, \dots, P_m that partition \mathbb{Z}_n and quadratic polynomials $q_1, \dots, q_m : \mathbb{Z}_n \rightarrow \mathbb{Z}_n$ such that*

$$n^{-1} \sum_i \left| \sum_{x \in P_i} f(x) \omega^{-q_i(x)} \right|$$

is bounded away from zero.

This tells us that on average f correlates with quadratic phase functions on the arithmetic progressions P_i . From this result it turns out to be possible to deduce that there is a refined partition into smaller arithmetic progressions such that f correlates on average with *linear*

phase functions, and then we are in essentially the situation we were in with Roth's theorem and can complete the proof of Szemerédi's theorem for progressions of length 4.

From the point of view of generalizing Fourier analysis, however, Lemma 10.2 is unsatisfactory. Our previous deductions tell us that the hypotheses of the lemma holds when f is a function with $\|f\|_\infty \leq 1$ and $\|f\|_{U^3} \geq c$, so the conclusions hold too. That gives us a lot of information about f , but it says nothing about how the quadratic polynomials q_i might be related. It therefore gives us only local information about f , from which it is not possible to deduce a converse: just because f correlates with quadratic phase functions on the progressions P_i , it does not follow that $\|f\|_{U^3}$ is large. (In fact, even constant functions do not do the job: if we were to choose for each i a random $\epsilon_i \in \{-1, 1\}$ and set $f(x)$ to equal ϵ_i everywhere on P_i , we would not have a function with large U^3 norm.)

By contrast, if $\|f\|_{U^2}$ is large, then we obtain very simply that $\|\hat{f}\|_\infty$ is large, which tells us that f correlates with a function of the form ω^{rx} , and that, equally simply, implies that $\|f\|_{U^2}$ is large.

What we would really like is to get from the hypothesis of Lemma 10.2 to a more global conclusion, that would say that f correlates with a generalized quadratic phase function of the kind described in the previous section. It is plausible that such a result should exist: from linear behaviour of the function ϕ one can deduce straightforwardly that f correlates with a pure quadratic phase function, so if we have generalized linear behaviour (of a rather precise kind) then it seems reasonable to speculate that f should correlate with a correspondingly generalized quadratic phase function.

The main obstacle to proving this is that the function $(a, x) \mapsto a\phi(x)$ is not symmetric. If it were, then the proof would be easy. However, Green and Tao found an ingenious "symmetrization argument" that allowed them to deduce from the hypotheses of Lemma 10.2 a more symmetric set of hypotheses that yielded the desired result. I shall state it somewhat imprecisely here. It is known as the inverse theorem for the U^3 norm.

Theorem 10.3. *For every $c > 0$ there exists $c' > 0$ with the following property. Let $f : \mathbb{Z}_n \rightarrow \mathbb{C}$ be a function with $\|f\|_\infty \leq 1$ and $\|f\|_{U^3} \geq c$. Then there exists a generalized quadratic phase function g such that $\langle f, g \rangle \geq c'$. Conversely, every function that correlates well with a generalized quadratic phase function has a large U^3 norm.*

The imprecision is of course that I have not said exactly what a quadratic phase function is. There are in fact several non-identical ways of defining them and the theorem is true for each such way. The way I presented them in the previous section (where the exponent is something like a quadratic form on a multidimensional arithmetic progression) is perhaps

the easiest to understand for a non-expert, but it is not the most convenient to use in proofs.

A natural question to ask at this point is what happens for the U^k norm when $k \geq 4$. If one is aiming for a generalization of Lemma 10.2, and thereby for a proof of Szemerédi's theorem, the case $k = 4$ (which corresponds to arithmetic progressions of length 5) is significantly harder than the case $k = 3$, and after that the difficulty does not increase further. As for the inverse theorem, one would like to show that a function with large U^k norm correlates well with a generalized polynomial phase function of degree $k - 1$, but it is far from easy even to come up with a satisfactory definition of what such a function should be. This Green and Tao did in a famous paper entitled Linear Equations in the Primes, which is about a programme to hugely extend their even more famous paper The Primes Contain Arbitrarily Long Arithmetic Progressions. Actually proving the resulting inverse conjecture took several years, but finally, with Tamar Ziegler, they managed it. So we now have something like a higher-order Fourier analysis for every degree.

11. HYPERGRAPHS

A graph is a collection of pairs of elements of a set. What happens if we generalize from pairs to triples and beyond? A k -uniform hypergraph is a set X and a subset of $X^{(k)}$, where $X^{(k)}$ denotes the set of all subsets of X of size k . In this section I shall concentrate on the case $k = 3$, though it should be fairly clear how to generalize what I say to higher values.

Just as it is natural, when one thinks about graphs in an analytic way, to think of them as special kinds of matrices, or functions of two variables, so hypergraphs can be thought of as functions of three variables. Furthermore, there is a natural three-variable analogue of the box norm that we saw earlier. It is given by the following formula.

$$\|f\|_{\square^3}^8 = \mathbb{E}_{x,x',y,y',z,z'} f(x,y,z) \overline{f(x,y,z')} \overline{f(x,y',z)} \overline{f(x,y',z')} \overline{f(x',y,z)} \overline{f(x',y,z')} \overline{f(x',y',z)} \overline{f(x',y',z')}.$$

As usual, one can define a corresponding box inner product by using eight different functions instead of just one, the inner product satisfies a Cauchy-Schwarz-type inequality, and that inequality can be used to prove that the norm really is a norm. Now let us look at some further useful facts about the box norm.

There is enough similarity between the formula for the box norm and the formula for the U^3 norm for it to be highly plausible that there should be a close relationship between

them. And indeed there is. Let G be a finite Abelian group, let $f : G \rightarrow \mathbb{C}$ be some function, and define a three-variable function $F : G^3 \rightarrow \mathbb{C}$ by $F(x, y, z) = f(x + y + z)$. It is easy to check directly from the formula that $\|F\|_{\square^3} = \|f\|_{U^3}$. (A similar relationship can also be shown between the two-dimensional box norm and the U^2 norm.)

It is a little surprising, therefore, that one can prove rather easily an inverse theorem for the box norm. As we shall see, however, the information it gives us is not strong enough to allow us to deduce from it the inverse theorem for the U^3 norm.

Let X be a finite set and let $f : X^3 \rightarrow \mathbb{C}$ be a function with $\|f\|_\infty \leq 1$ and $\|f\|_{\square^3} \geq c$. The second inequality tells us that

$$\mathbb{E}_{x,x',y,y',z,z'} f(x, y, z) \overline{f(x, y, z')} \overline{f(x, y', z)} \overline{f(x, y', z')} \\ \overline{f(x', y, z)} \overline{f(x', y, z')} \overline{f(x', y', z)} \overline{f(x', y', z')} \geq c^8.$$

By averaging, there must exist x', y', z' such that

$$|\mathbb{E}_{x,y,z} f(x, y, z) \overline{f(x, y, z')} \overline{f(x, y', z)} \overline{f(x, y', z')} \\ \overline{f(x', y, z)} \overline{f(x', y, z')} \overline{f(x', y', z)} \overline{f(x', y', z')}| \geq c^8.$$

We can think of the left-hand side as the modulus of the inner product of f with the function g , given by the formula

$$g(x, y, z) = f(x, y, z') \overline{f(x, y', z)} \overline{f(x, y', z')} \overline{f(x', y, z)} \overline{f(x', y, z')} \overline{f(x', y', z)} \overline{f(x', y', z')}.$$

The interesting thing about g is that it is a product of functions each of which depends on at most two of the variables x, y, z . Thus, we find that if $\|f\|_{\square^3}$ is large, then it correlates with a function of “lower complexity”. The analogue of these low-complexity functions for matrices is the matrices of the form $u \otimes v$ – that is, the matrices of rank 1.

However, in the two-variable case we have more. For Hermitian matrices we have a decomposition of the form $\sum_i \lambda_i u_i \otimes \overline{u_i}$, where (u_i) is an orthonormal basis, and in general we have a singular-value decomposition $\sum_i \lambda_i u_i \otimes \overline{v_i}$, where (u_i) and (v_i) are both orthonormal. If u, v and w are functions of two variables, let us write $[[u, v, w]]$ for the function whose value at (x, y, z) is $u(x, y)v(y, z)w(z, x)$. Then the very simple inverse theorem just proved tells us that a function with large box norm correlates with a function of the form $[[u, v, w]]$, but what we do not seem to have is a canonical way of decomposing an arbitrary function as a sum of the form $\sum_i \lambda_i [[u_i, v_i, w_i]]$.

What happens if we try to deduce the inverse theorem for the U^3 norm from the inverse theorem for the box norm in three variables? If $\|f\|_{U^3} \geq c$, then the argument gives us functions f_1, \dots, f_6 , all of ℓ_∞ norm at most 1, such that

$$|\mathbb{E}_{x,y,z} f(x+y+z) \overline{f_1(x+y) f_2(y+z) f_3(z+x)} f_4(x) f_5(y) f_6(z)| \geq c^8.$$

However, it does not tell us anything much about the structure of the functions f_1, \dots, f_6 . It is possible to deduce from the inequality above that they have quadratic structure, and that the inverse theorem therefore holds, but the proof is no easier than the proof of the inverse theorem was already – it just uses the same general approach in an unnecessarily complicated way.

Despite this, the theory of hypergraphs has been important and useful in additive combinatorics. I will not explain why here, except to mention a theorem about hypergraphs that turns out to imply Szemerédi’s theorem, known as the *simplex removal lemma*. (The implication is fairly straightforward, but slightly too long to give here.) Define a *simplex* in a k -uniform hypergraph H to be a set of $k+1$ vertices such that any k of them form an edge H . (The word “edge” here means one of the sets of size k that belongs to H . When $k=2$, a simplex is a triangle.)

Theorem 11.1. *For every $c > 0$ and positive integer k there exists $a > 0$ with the following property. If H is a k -uniform hypergraph with n vertices that contains at most an^{k+1} simplices, then it is possible to remove at most cn^k edges from H to create a k -uniform hypergraph that contains no simplices at all.*

Rather surprisingly, even when $k=2$, when the result says that a graph with few triangles is close to a graph with no triangles, this result is not straightforward. In particular, the best known dependence of a on c is extremely weak: its reciprocal is a tower of 2s of height proportional to $\log(1/c)$. So a bound of the form $\exp(-(1/c)^A)$ for some fixed $A > 0$ would, for example, be a major improvement.

12. FOURIER ANALYSIS ON NON-ABELIAN GROUPS

The following result is easy to prove. We say that a subset of an Abelian group is *sum free* if it contains no three elements x, y, z with $x + y = z$.

Theorem 12.1. *There exists a constant $c > 0$ such that every finite Abelian group G has a subset A of cardinality at least $c|G|$ that is sum free.*

To see this, let \mathbb{Z}_m be one of the cyclic groups of which G is a product, and take all elements whose coordinate in this copy of \mathbb{Z}_m lies between $m/3$ and $2m/3$ (and strictly between on one of the two sides).

Babai and Sós asked whether a similar result held for general finite groups. They almost certainly expected the answer no, but it turns out not to be completely obvious how to disprove it.

One thing it is natural to do is to look at groups that are “highly non-Abelian”. This can be measured in various ways. One is to look at the sizes of conjugacy classes. If a group G is Abelian, then all its conjugacy classes are singletons, so if a group has large conjugacy classes, then that is saying that in some sense it is far from Abelian: not only are the conjugates gxg^{-1} not all equal to x , they are not even concentrated in a small subset of the group.

Another property that characterizes Abelian groups is that all their irreducible representations are one-dimensional. So another potential way of measuring non-Abelianness is to look at the lowest dimension of an irreducible representation.

Since we have already made use of characters of finite Abelian groups – that is, their irreducible representations – and since we are trying to count solutions to a simple equation in a dense subset of a group, the second measure looks promising. And it does indeed turn out to be possible to solve this problem by using a more general Fourier analysis, in which characters are replaced by more general irreducible representations.

The definition of the Fourier transform of a function $f : G \rightarrow \mathbb{C}$ is more or less the first thing one writes down. If $\rho : G \rightarrow U(k)$ is an irreducible unitary representation of G , then

$$\hat{f}(\rho) = \mathbb{E}_x f(x) \overline{\rho(x)}.$$

(Another candidate for the definition would be as above but with the conjugate $\overline{\rho(x)}$ replaced by the adjoint $\rho(x)^*$, but the conjugate turns out to be more convenient.)

For this to be a useful definition, we would like it to satisfy natural analogues of the basic properties of the Abelian Fourier transform. And indeed it does. Parseval’s identity, for example, takes the following form. If f and g are functions from G to \mathbb{C} , then

$$\mathbb{E}_x f(x) \overline{g(x)} = \sum_{\rho} n_{\rho} \text{tr}(\hat{f}(\rho) \hat{g}(\rho)^*),$$

where the sum is over all irreducible representations and for each such representation ρ its dimension is n_ρ . Let us briefly see how this is proved. We have

$$\begin{aligned} \sum_{\rho} n_{\rho} \operatorname{tr}(\hat{f}(\rho) \hat{g}(\rho)^*) &= \sum_{\rho} n_{\rho} \mathbb{E}_{x,y} f(x) \overline{g(y)} \operatorname{tr}(\rho(x) \rho(y)^*) \\ &= \mathbb{E}_{x,y} f(x) \overline{g(y)} \sum_{\rho} n_{\rho} \overline{\operatorname{tr}(\rho(x) \rho(y)^*)} \end{aligned}$$

We now use a fundamental orthogonality result from basic representation theory, which states that $\sum_{\rho} n_{\rho} \operatorname{tr}(\rho(x) \rho(y)^*) = n$ if $x = y$ and 0 otherwise. It follows that

$$\mathbb{E}_{x,y} f(x) \overline{g(y)} \sum_{\rho} n_{\rho} \overline{\operatorname{tr}(\rho(x) \rho(y)^*)} = \mathbb{E}_x f(x) \overline{g(x)}$$

and the proof is complete.

How about the convolution identity? It states, as we would hope, that

$$\widehat{f * g}(\rho) = \hat{f}(\rho) \hat{g}(\rho)$$

for any two functions $f, g : G \rightarrow \mathbb{C}$ and any irreducible representation ρ . Again it is instructive to see the proof. We have

$$\begin{aligned} \widehat{f * g}(\rho) &= \mathbb{E}_x (f * g)(x) \overline{\rho(x)} \\ &= \mathbb{E}_x \mathbb{E}_{uv=x} f(u) g(v) \overline{\rho(x)} \\ &= \mathbb{E}_{u,v} f(u) g(v) \overline{\rho(u) \rho(v)} \\ &= (\mathbb{E}_u f(u) \overline{\rho(u)}) (\mathbb{E}_v g(v) \overline{\rho(v)}) \\ &= \hat{f}(\rho) \hat{g}(\rho). \end{aligned}$$

Note that we used the fact that $\overline{\rho(uv)} = \overline{\rho(u) \rho(v)}$ in the proof above. Had we defined the Fourier transform using adjoints, we would have had to use instead the fact that $\rho(uv)^* = \rho(v)^* \rho(u)^*$, so we would have obtained the identity $\widehat{f * g}(\rho) = \hat{g}(\rho) \hat{f}(\rho)$.

The last property I want to discuss is the inversion formula. Here we have what looks at first like a puzzle: in the Abelian case we decomposed functions as linear combinations of characters, but irreducible representations are matrix-valued functions of different dimensions, so we cannot express scalar-valued functions as linear combinations of them.

There is of course a natural way of converting a matrix-valued function into a scalar-valued function, and that is to take the trace. Moreover, traces of representations are well known to be important functions – they are characters in the sense of representation theory.

So can we decompose a function as a linear combination of functions of the form $\chi(x) = \text{tr}(\rho(x))$? No we cannot, since such functions are constant on conjugacy classes. (We can, however, decompose functions if they *are* constant on conjugacy classes – such functions are called *class functions*.) In fact, since there are not n inequivalent irreducible representations (except when the group is Abelian), there is no hope of writing down some scalar-valued functions u_ρ and expanding every f as a linear combination of the u_ρ .

However, we shouldn't necessarily expect to be able to do so. We would like the coefficients in our inversion formula to be the matrices $\hat{f}(\rho)$ in some suitable sense. And once we make that our aim, it is a short step to writing down the following slightly subtler formula.

$$f(x) = \sum_{\rho} n_{\rho} \text{tr}(\hat{f}(\rho) \overline{\rho(x)^*}).$$

This can be verified easily using the orthogonality property we used earlier.

It is not hard to check that this formula specializes to the formula given earlier when the group is Abelian. One way of making it look more like that formula is to define \hat{G} to be the set of all irreducible representations of G (up to equivalence), to define $M(\hat{G})$ to be the set of all matrix-valued functions \hat{f} on \hat{G} such that $\hat{f}(\rho)$ is an $n_{\rho} \times n_{\rho}$ matrix for every ρ , and to define an inner product on $M(\hat{G})$ by the formula

$$\langle \hat{f}, \hat{g} \rangle = \sum_{\rho} n_{\rho} \langle \hat{f}(\rho), \hat{g}(\rho) \rangle,$$

where the inner product on the right-hand side is the matrix inner product $\langle A, B \rangle = \text{tr}(AB^*) = \sum_{i,j} A_{ij} \overline{B_{ij}}$. Note that Parseval's identity now becomes the usual formula $\langle f, g \rangle = \langle \hat{f}, \hat{g} \rangle$. As for the inversion formula, it can be written as follows.

$$f(x) = \langle \hat{f}, \delta_x^* \rangle,$$

where δ_x^* is the evaluation function $\rho \mapsto \rho(x)^*$. The right-hand side can be expanded to $\sum_{\rho} n_{\rho} \langle \hat{f}(\rho), \overline{\rho(x)} \rangle$, which is equal to $\sum_{\chi} \hat{f}(\chi) \chi(x)$ when G is Abelian.

Before we explain how the problem of Babai and Sós can be solved, we need a couple of simple results about matrices. I should remark that because we are on the Fourier side, we are dealing with sums rather than expectations. In particular, we are using the standard notion of matrix multiplication, and the box norm will be defined using sums.

Lemma 12.2. *Let A and B be square matrices. Then $\|AB\|_{HS} \leq \|A\|_{\square} \|B\|_{\square}$.*

Proof. Observe that

$$\begin{aligned} \|AB\|_{HS}^2 &= \sum_{x,x'} \left| \sum_y A(x,y)B(y,x') \right|^2 \\ &= \sum_{x,x'} \sum_{y,y'} A(x,y) \overline{A(x,y')} \overline{B^*(x',y)} B^*(x',y'). \end{aligned}$$

This last expression is the ‘‘box inner product’’ $[A, A, B^*, B^*]$, and as we saw earlier (in the section on matrices) it satisfies a Cauchy-Schwarz-type inequality

$$[A, B, C, D] \leq \|A\|_{\square} \|B\|_{\square} \|C\|_{\square} \|D\|_{\square}.$$

Applying this, together with the fact that $\|B^*\|_{\square} = \|B\|_{\square}$, we obtain the result. \square

Lemma 12.3. *For every matrix A we have $\|A\|_{\square} \leq \|A\|_{HS}$.*

Proof. This can be shown with a direct argument, but it also follows from the fact that $\|A\|_{\square}$ is the ℓ_4 norm of the singular values of A and $\|A\|_{HS}$ is the ℓ_2 norm. \square

Lemma 12.4. *Let G be a finite group and let $f, g : G \rightarrow \mathbb{C}$ be functions with average zero. Let m be the smallest dimension of a non-trivial representation of G . Then*

$$\|f * g\|_2 \leq m^{-1/2} \|f\|_2 \|g\|_2.$$

Proof. By the convolution identity, Parseval’s identity and the lemmas above, we have that

$$\begin{aligned} \|f * g\|_2^2 &= \sum_{\rho} n_{\rho} \|\hat{f} \hat{g}\|_{HS}^2 \\ &\leq \sum_{\rho} n_{\rho} \|\hat{f}\|_{\square}^2 \|\hat{g}\|_{\square}^2 \\ &\leq \sum_{\rho} n_{\rho} \|\hat{f}\|_{HS}^2 \|\hat{g}\|_{HS}^2. \end{aligned}$$

Since f averages zero, $\hat{f}(\rho) = 0$ when ρ is the trivial representation. Also, by Parseval’s identity we have that $\sum_{\rho} n_{\rho} \|\hat{f}(\rho)\|_{HS}^2 = \|f\|_2^2$. It follows that the maximum possible value of $\|\hat{f}(\rho)\|_{HS}^2$ is $m^{-1} \|f\|_2^2$. Therefore, using Parseval’s identity again, we find that

$$\sum_{\rho} n_{\rho} \|\hat{f}\|_{HS}^2 \|\hat{g}\|_{HS}^2 \leq m^{-1} \|f\|_2^2 \sum_{\rho} n_{\rho} \|\hat{g}(\rho)\|_{HS}^2 = m^{-1} \|f\|_2^2 \|g\|_2^2,$$

which completes the proof. \square

Now let us quickly deduce that if a group G has no non-trivial low-dimensional representations, then it does not contain a large product-free set.

Theorem 12.5. *Let G be a finite group and let m be the smallest dimension of a non-trivial representation of G . Then G contains no product-free subset of density greater than $m^{-1/3}$.*

Proof. Let α be the density of A and as usual let f be the function $f(x) = A(x) - \alpha$. We shall now try to show that

$$\mathbb{E}_{xy=z} A(x)A(y)A(z) \neq 0,$$

which obviously implies that A is not product free.

We have that

$$\mathbb{E}_{xy=z} A(x)A(y)A(z) = \mathbb{E}_{xy=z} (\alpha + f(x))(\alpha + f(y))(\alpha + f(z)),$$

and since f averages zero, if we expand the right-hand side into eight separate sums, we find that all terms are zero apart from two, and we obtain the expression

$$\alpha^3 + \mathbb{E}_{xy=z} f(x)f(y)f(z) = \alpha^3 + \langle f * f, f \rangle.$$

By the Cauchy-Schwarz inequality and Lemma 12.4 we have that

$$|\langle f * f, f \rangle| \leq \|f * f\|_2 \|f\|_2 \leq m^{-1/2} \|f\|_2^3.$$

We also have that $\|f\|_2^2 = \alpha(1 - \alpha)^2 + (1 - \alpha)\alpha^2 = \alpha(1 - \alpha)$. Therefore, if A is product free we must have the inequality

$$\alpha^{3/2}(1 - \alpha)^{3/2}m^{-1/2} \geq \alpha^3,$$

which implies that $\alpha \leq m^{-1/3}$. \square

It remains to remark that there do exist groups with no low-dimensional representations. Indeed, any family of finite simple groups has this property, though some have it much more strongly than others. The “most” non-Abelian family of groups is the family $\text{PSL}(2, q)$. If a group in this family has order n , then its non-trivial representations have dimension at least $cn^{1/3}$, where $c > 0$ is an absolute constant. Therefore, these groups have no product-free subsets of density greater than $c'n^{-1/9}$.

The same argument shows that if A , B and C are sets of density greater than $m^{-1/3}$, then $ABC = \{abc : a \in A, b \in B, c \in C\} = G$.

There turns out to be a close connection between groups with no low-dimensional representations and quasirandom graphs. If G is a finite group with no low-dimensional non-trivial representations, then for any dense set $A \subset G$ we can define a bipartite graph with two copies of G as its vertex sets and x joined to y if and only if $y = ax$ for some $a \in A$. The remark above about sets A, B and C tells us that this graph is quasirandom.

As a final remark, we note that the U^2 norm can be generalized easily to a non-Abelian context. A good definition turns out to be as follows.

$$\|f\|_{U^2}^4 = \mathbb{E}_{xy^{-1}zw^{-1}=e} f(x)\overline{f(y)}f(z)\overline{f(w)}.$$

The properties of this norm are just what one would hope. For example, one can define a generalized inner product in the obvious way, and we do indeed have the inequality

$$[f_1, f_2, f_3, f_4] \leq \|f_1\|_{U^2}\|f_2\|_{U^2}\|f_3\|_{U^2}\|f_4\|_{U^2},$$

which can then be used in the usual way to prove that this U^2 norm is a norm. We also have the Fourier interpretation that one would guess, namely

$$\|f\|_{U^2}^4 = \sum_{\rho} n_{\rho} \|\hat{f}(\rho)\|_{\square}^4.$$

This is the natural guess because it involves fourth powers on the right-hand side, both in the obvious sense that there is a fourth power visible in the expression, and also in the less obvious sense that the box norm of a matrix is equal to the ℓ_4 norm of the singular values. (Thus, in a certain sense we have fourth powers of generalized Fourier coefficients in two different ways.) Indeed, there is a natural way of defining an ℓ_p norm on \hat{G} for every p . For an $m \times m$ matrix A , one defines the *trace-class norm* $\|A\|_p$ to be the ℓ_p norm of the singular values of A , and then for a matrix-valued function \hat{f} one defines $\|\hat{f}\|_p$ by the formula

$$\|\hat{f}\|_p^p = \sum_{\rho} n_{\rho} \|\hat{f}(\rho)\|_p^p.$$

That is, we take the ℓ_p norm in \hat{G} of the function $x \mapsto \|\hat{f}(x)\|_p$. Once we have done this, we have the familiar identity

$$\|f\|_{U^2} = \|\hat{f}\|_4.$$

Functions with small U^2 norms (given their averages) behave like random functions, and when a group has no non-trivial low-dimensional representations, Lemma 12.4 tells us that *all* reasonably spread out functions behave like random functions. To see this, note that

$\|f\|_{U^2}^4 = \|f * f^*\|_2^2$, where $f^*(x)$ is defined to be $\overline{f(x^{-1})}$, so if f averages zero, then we have the inequality $\|f\|_{U^2}^4 \leq m^{-1}\|f\|_2^2$.

13. FOURIER ANALYSIS FOR MATRIX-VALUED FUNCTIONS

Let G be a finite group and let $f : G \rightarrow M_n(\mathbb{C})$ be a matrix-valued function. (We are not assuming that n is the order of G .) We can define a Fourier transform for f by simply applying the definition of the previous section to each matrix coefficient. That is, for each $i, j \leq n$ we define f_{ij} to be the function $x \mapsto f(x)_{ij}$, and then for each irreducible representation ρ we define $\hat{f}(\rho)$ to be the $n \times n$ block matrix whose ij th entry is the $n_\rho \times n_\rho$ matrix $\widehat{f_{ij}}(\rho)$. Thus, $\hat{f}(\rho)$ is an $nn_\rho \times nn_\rho$ matrix.

We can write this definition more concisely as follows.

$$\hat{f}(\rho) = \mathbb{E}_x f(x) \otimes \overline{\rho(x)}.$$

Thus, whereas with Abelian groups we had scalar-valued functions and scalar-valued representations, and in the previous section we had scalar-valued functions and matrix-valued representations, now we have matrix-valued functions and matrix-valued representations. In each case we take tensor products, but in the first two cases they are trivial.

In order to state the basic properties of the Fourier transform, we need to be clear about our notation. For a matrix-valued function f on the physical side, we shall write $\|f\|_2$ for the norm defined by the formula

$$\|f\|_2^2 = \mathbb{E}_x \|f(x)\|_2^2$$

Here it turns out to be convenient to take $\|A\|_2$ to be the non-normalized Hilbert-Schmidt norm. Thus, the norm scales with the dimension of f , but not with the size of the group.

On the Fourier side, we have similar definitions but using sums all the way through, so these are the same as the definitions of the norms in the scalar case.

With these normalizations, the first few basic properties of the Fourier transform now read as follows.

- $\|f\|_2^2 = \|\hat{f}\|_2^2$. (Parseval's identity).
- $\langle f, g \rangle = \langle \hat{f}, \hat{g} \rangle$. (Parseval's identity).
- $\widehat{f * g}(\rho) = \hat{f}(\rho)\hat{g}(\rho)$. (Convolution formula).

The proofs are more or less the same as in the scalar case, but to clarify the point about normalizations, we give a proof of the second version of Parseval's identity, which goes like this. (It is important to realize that the meaning of the inner product varies from expression

to expression – sometimes we are talking about the inner product of two matrices, and sometimes about the inner product of two matrix-valued functions.)

$$\begin{aligned} \langle \hat{f}, \hat{g} \rangle &= \sum_{\rho} n_{\rho} \langle \hat{f}(\rho), \hat{g}(\rho) \rangle \\ &= \sum_{\rho} n_{\rho} \langle \mathbb{E}_x f(x) \otimes \overline{\rho(x)}, \mathbb{E}_y g(y) \otimes \overline{\rho(y)} \rangle \\ &= \mathbb{E}_{x,y} \langle f(x), g(y) \rangle \sum_{\rho} n_{\rho} \langle \overline{\rho(x)}, \overline{\rho(y)} \rangle \end{aligned}$$

In the last expression, both inner products use sums.

By the basic orthogonality property from representation theory, the sum over ρ is equal to $|G|\delta_{xy}$, so we end up with $\mathbb{E}_x \langle f(x), g(x) \rangle$, which is the definition of $\langle f, g \rangle$.

The inversion formula is also straightforward, but it needs a little notation. Recall that for scalar-valued functions the inversion formula was $f(x) = \sum_{\rho} n_{\rho} \text{tr}(\hat{f}(\rho) \overline{\rho(x)^*})$. Since the Fourier transform for matrix-valued functions is obtained by applying the Fourier transform for scalar-valued functions to each matrix entry, we obtain the formula

$$f(x)_{ij} = \sum_{\rho} n_{\rho} \text{tr}(\widehat{f}_{ij}(\rho) \overline{\rho(x)^*}).$$

Let us write tr_{ρ} for the operation that takes an $n \times n$ block matrix A with blocks that are $n_{\rho} \times n_{\rho}$ matrices and returns the $n \times n$ matrix whose ij th value is the (unnormalized) trace of the ij th block of A . Then we can write the inversion formula in the form

$$f(x) = \sum_{\rho} \text{tr}_{\rho}(\hat{f}(\rho) \overline{\rho(x)^*}),$$

which is just like the formula when f takes scalar values except that the trace function tr has been replaced by the matrix-of-traces function tr_{ρ} .

What is this definition good for? At least one situation in which it is useful is when one wishes to measure the extent to which a matrix-valued function behaves like a representation. To illustrate this, let me discuss a very nice result of Moore and Russell (not the famous philosophers, but a pair of contemporary mathematicians).

In the previous section, we remarked that if G is a finite group with no low-dimensional non-trivial representations, then for every function $f : G \rightarrow \mathbb{C}$ such that $\|f\|_{\infty} \leq 1$ and

$\mathbb{E}_x f(x) = 0$ the U^2 norm of f is small. That is,

$$\mathbb{E}_{xy^{-1}zw^{-1}} f(x) \overline{f(y)} \overline{f(z)} \overline{f(w)}$$

is small. Now if we could find a non-trivial character of G , in the sense of a homomorphism from G to \mathbb{C} that is not the identity, then this would not be true: whenever $xy^{-1}zw^{-1} = e$ we would have $f(x) \overline{f(y)} \overline{f(z)} \overline{f(w)} = 1$ so the average would be 1, which is the largest it can possibly be. So the observation that the U^2 norm has to be small is telling us that G not only fails to have a non-trivial character (which we know because it has no non-trivial low-dimensional representations), but it does not admit any functions that are even very slightly *close* to being a non-trivial character: if f averages zero, then it not possible for the average real part of $f(x) \overline{f(y)} \overline{f(z)} \overline{f(w)}$ to be greater than some small constant when $xy^{-1}zw^{-1} = e$.

Moore and Russell showed that this observation can be extended to matrix-valued functions. (Actually, the precise result they showed was not quite this one, but it was very similar and can be proved by similar methods.) That is, if the dimension of the smallest non-trivial representation is m , if n is substantially less than m , and if f takes values that are $n \times n$ matrices with operator norm at most 1, and if $\mathbb{E}_x f(x) = 0$, then the U^2 norm of f is small. This is a stronger statement, because matrix-valued functions have more elbow room and therefore more room to create the necessary correlations. It is also stronger in a more obvious way: it tells us that not only are scalar-valued functions on G as unlike non-trivial characters as they could possibly be, low-dimensional matrix-valued functions are as far from non-trivial representations as they could possibly be.

We begin by observing that the statement and proof of Lemma 12.4 carry over almost word for word to the matrix-valued case. (The proof is identical, so we do not give it again.)

Lemma 13.1. *Let G be a finite group and let m be the smallest dimension of a non-trivial representation of G . Let $f, g : G \rightarrow M_n(\mathbb{C})$ be two matrix-valued functions that average zero. Then*

$$\|f * g\|_2 \leq m^{-1/2} \|f\|_2 \|g\|_2.$$

There is, however, an important new factor to take into account here, which is that the two sides of the inequality scale differently with the dimension. Suppose, for instance, that both f and g are equal to the same n -dimensional representation ρ . Then the 2-norms of every single $f(x)$, $g(x)$ and $f * g(x)$ are all equal to n , so $\|f * g\|_2^2 = n$, while $m^{-1} \|f\|_2^2 \|g\|_2^2 = m^{-1} n^2$. So the inequality does not stop f and g from being representations

when $n = m$, which of course is as it should be, since m is defined to be the dimension of a representation of G .

With that remark in mind, let us turn to U^2 norms. The natural definition of the U^2 norm in the matrix-valued case is

$$\|f\|_{U^2}^4 = \mathbb{E}_{xy^{-1}zw^{-1}=e} \operatorname{tr}(f(x)f(y)^*f(z)f(w)^*),$$

where the trace is not normalized. This is equal to $\|f * f^*\|_2^2$ and it is also equal to $\|\hat{f}\|_4^4$. (Recall that this is defined to be $\sum_{\rho} n_{\rho} \|\hat{f}\|_{\square}^4$.) Therefore, by Lemma 13.1 we find that

$$\|f\|_{U^2}^4 \leq m^{-1} \|f\|_2^4.$$

In particular, if f takes values $f(x)$ with Hilbert-Schmidt norm at most $n^{1/2}$ (which is the case, for example, if they are all unitary, and more generally if they all have operator norm at most 1), then $\|f\|_{U^2}^4 \leq m^{-1} n^2$. For f to be a unitary representation, we would need $\|f\|_{U^2}^4$ to be equal to n (since $f(x)f(y)^*f(z)f(w)^*$ would be the identity whenever $xy^{-1}zw^{-1} = e$), which would equal $n^{-1} \|f\|_2^4$, which is at most n , by hypothesis. Therefore, if n is significantly less than m , so that $m^{-1} n^2$ is significantly less than n , we see that f is not even close to being a representation, in the sense that there is almost no correlation between $f(x)f(y)^*$ and $f(w)f(z)^*$ even if we are given that $xy^{-1} = wz^{-1}$.

14. AN INVERSE THEOREM FOR THE MATRIX U^2 NORM

Recall the very simple inequalities that we used earlier to relate the ℓ_4 and ℓ_{∞} norms of the Fourier transform of a function f from an Abelian group to \mathbb{C} . If we know that $\|f\|_2 \leq 1$, then we find that

$$\|\hat{f}\|_{\infty}^4 \leq \|\hat{f}\|_4^4 \leq \|\hat{f}\|_{\infty}^2 \|\hat{f}\|_2^2 \leq \|\hat{f}\|_{\infty}^2.$$

Since $\|\hat{f}\|_4 = \|\hat{f}\|_{U^2}$, we deduce that if $\|f\|_{U^2} \geq c$, then there exists a character χ such that $|\langle f, \chi \rangle| = |\hat{f}(\chi)| \geq c^2$.

What happens if we try to generalize this to non-Abelian groups and to matrix-valued functions? Let us assume that $f(x)$ is an $n \times n$ matrix with operator norm at most 1 for every x . (The operator norm is the maximum of the singular values, and thus the natural ℓ_{∞} norm of a matrix.) Then just as before, we have

$$\|\hat{f}\|_{\infty}^4 \leq \|\hat{f}\|_4^4 = \sum_{\rho} n_{\rho} \|\hat{f}(\rho)\|_4^4 \leq \max_{\rho} \|\hat{f}(\rho)\|_{\infty}^2 \sum_{\rho} n_{\rho} \|\hat{f}(\rho)\|_2^2 = \|\hat{f}\|_{\infty}^2 \|f\|_2^2.$$

Unfortunately, if n is large, then this is no longer a rough equivalence, since the best we can say about $\|f\|_2^2$ is that it is at most n (since it is the sum of the squares of n singular values, each of which lies between 0 and 1).

However, that does not mean that there is nothing we can say. The largest possible value of $\|\hat{f}\|_4^4 = \|f\|_{U^2}^4$ is, as we have seen, n . Let us take a function f such that $\|f\|_{U^2}^4 \geq cn$. Then we obtain from the second inequality above that there exists an irreducible representation ρ such that $\|\hat{f}(\rho)\|_\infty \geq c^{1/2}$.

If G is Abelian and $f : G \rightarrow \mathbb{C}$, then $\|\hat{f}(\chi)\|_\infty$ is just $|\langle f, \chi \rangle|$, so this statement is saying that f correlates in a significant way with a character. But in our situation we have the more complicated statement that

$$\|\mathbb{E}_x f(x) \otimes \overline{\rho(x)}\|_\infty \geq c^{1/2}.$$

Is this telling us that f correlates in some sense with ρ ?

Let us try to interpret it. We shall first use the fact that $\|A\|_\infty$, the operator norm of A , is the largest possible value of $\|Au\|_2$ over all unit vectors u , which in turn is the largest possible value of $\langle Au, v \rangle$ over all pairs of unit vectors u and v . Therefore, we can find unit vectors u and v such that

$$\langle (\mathbb{E}_x f(x) \otimes \overline{\rho(x)})u, v \rangle \geq c^{1/2}.$$

Now let us rewrite this in coordinate form. Because of the special form of the $n \times n_\rho$ matrix $f(x) \otimes \overline{\rho(x)}$ it is natural to give it four indices instead of two: we have that $(f(x) \otimes \overline{\rho(x)})_{ijkl} = f(x)_{ik} \overline{\rho(x)_{jl}}$. Then indexing u and v in the corresponding way, and writing them as U and V since they have now become matrices, we have that

$$\langle (f(x) \otimes \overline{\rho(x)})U, V \rangle = \sum_{i,j,k,l} f(x)_{ik} \overline{\rho(x)_{jl}} U_{kl} \overline{V_{ij}} = \langle f(x)U \rho(x)^*, V \rangle,$$

where the product in the last expression is just normal matrix multiplication. Taking expectations, we deduce that

$$\mathbb{E}_x \langle f(x)U \rho(x)^*, V \rangle \geq c^{1/2}$$

for two $n \times n_\rho$ matrices U and V that have Hilbert-Schmidt norm 1. We can write this more symmetrically as

$$\mathbb{E}_x \langle f(x)U, V \rho(x) \rangle \geq c^{1/2}.$$

It is natural to rescale U and V so that they have Hilbert-Schmidt norm $n_\rho^{1/2}$. That is, we can say that there exist an irreducible representation ρ and matrices U and V with

$\|U\|_2^2 = \|V\|_2^2 = n_\rho$ such that

$$\mathbb{E}_x \langle f(x)U\rho(x)^*, V \rangle = \mathbb{E}_x \langle f(x)U, V\rho(x) \rangle = \mathbb{E}_x \langle f(x), V\rho(x)U^* \rangle \geq c^{1/2}n_\rho.$$

This seems quite satisfactory, but it falls short of being a true inverse theorem for the matrix U^2 norm because the converse does not hold. That is, if we are given ρ, U and V satisfying the above conditions, we cannot deduce that f has a large U^2 norm unless n_ρ is comparable to n , which it does not have to be.

However, we have not yet exhausted our options. If $\|f\|_{U^2}^4 = \|\hat{f}\|_4^4 \geq cn$, then we are given that

$$\sum_{\rho} n_{\rho} \|\hat{f}(\rho)\|_4^4 \geq cn,$$

where $\|\hat{f}(\rho)\|_4^4$ denotes the sum of the fourth powers of the singular values of $\hat{f}(\rho)$. Let these singular values be $\lambda_{\rho,i}$ for $i = 1, 2, \dots, n_\rho$. Then we find that

$$\sum_{\rho} n_{\rho} \sum_{i=1}^{n_{\rho}} \lambda_{\rho,i}^4 \geq cn.$$

From Parseval's inequality and the assumption that each $f(x)$ has operator norm at most 1 (and hence Hilbert-Schmidt norm at most n) we also have that

$$\sum_{\rho} n_{\rho} \sum_{i=1}^{n_{\rho}} \lambda_{\rho,i}^2 \leq n.$$

Also, since $\hat{f}(\rho) = \mathbb{E}_x f(x) \otimes \overline{\rho(x)}$ is an average of matrices with operator norm at most 1, every $\lambda_{\rho,i}$ is at most 1.

Let $\lambda_1, \dots, \lambda_m$ be the singular values $\lambda_{\rho,i}$ arranged in some order, and for each i let n_i be the n_ρ that corresponds to λ_i . Then we can rewrite these inequalities as

$$\begin{aligned} \sum_i n_i \lambda_i^4 &\geq cn, \\ \sum_i n_i \lambda_i^2 &\leq n, \end{aligned}$$

and

$$\lambda_i \leq 1.$$

Note that if $c = 1$, then the only way of achieving the above inequalities is for λ_i^4 to equal λ_i^2 for every i (assuming that none of the n_i is zero). Thus, each λ_i is either 0 or 1, and $\sum \{n_i : \lambda_i = 1\} = n$. It is a straightforward exercise to prove that the more

relaxed assumptions above lead to similar but more relaxed conclusion: we can find a set A and constants $c_1 > 0$ and C that depend on c only (with a power dependence) such that $c_1 n \leq \sum_{i \in A} n_i \leq Cn$, and $\lambda_i \geq c_1$ for every $i \in A$. In short, we can find a set of large singular values (coming from the various $\hat{f}(\rho)$ of size roughly comparable to n).

With each such singular value λ_i we can associate $n \times n_i$ matrices U_i and V_i with Hilbert-Schmidt norm n_i such that $\mathbb{E}_x \langle f(x)U_i, V_i \rho_i \rangle \geq \lambda_i n_i$, where ρ_i is the representation corresponding to λ_i . Moreover, if two pairs (U_i, V_i) and (U_j, V_j) come from the same ρ , then because of the nature of singular value decompositions, we have that $\langle U_i, U_j \rangle = \langle V_i, V_j \rangle = 0$.

I hope I have now written enough to make it plausible that we can put together these matrices and irreducible representations to create a representation-like function that correlates with f , and moreover that gives us an inverse theorem in the sense that the correlation in its turn implies that f has a large U^2 norm. Exactly how the putting together should work is not obvious, but it turns out that it can be done, and that it yields the following theorem. In the statement, recall that $\|f\|_\infty$ means the largest operator norm of any $f(x)$. Also, we define a *partial unitary* matrix to be an $n \times m$ matrix such that the rows are orthonormal if $n \leq m$ and the columns are orthonormal if $n \geq m$. (In particular, if $n = m$ then the matrix is unitary.)

Theorem 14.1. *Let G be a finite group, let $c > 0$ and let $f : G \rightarrow M_n(\mathbb{C})$ be a function such that $\|f\|_\infty \leq 1$ and $\|f\|_{U^2}^4 \geq cn$. Then there exists m such that $cn/4 \leq m \leq 4n/c$, an m -dimensional representation σ , and $n \times m$ partial unitary matrices U and V , such that*

$$\mathbb{E}_x \langle f(x)U, V\sigma(x) \rangle \geq c^2 m/16.$$

Note that σ will *not* normally be irreducible. This theorem tells us that f correlates with the function $V\sigma U^*$. The extra strength of this theorem over what we remarked earlier is that the dimension of σ is comparable to that of f . It turns out to be simple to deduce a converse statement – i.e., that if f correlates with a function of the above form, then $\|f\|_{U^2}^4 \geq c'n$ for some suitable c' . Therefore, Theorem 14.1 is indeed an inverse theorem for the matrix U^2 norm.

It is possible to give a more careful argument when $c = 1 - \epsilon$ for some small ϵ that allows us to show that $(1 - 2\epsilon)n \leq m \leq (1 - 4\epsilon)^{-1}n$, and to obtain a lower bound of $(1 - 16\epsilon)m$ in the last inequality. From this result it is not too hard to deduce a so-called *stability theorem* for near representations. Roughly speaking, it states that any unitary-valued function that almost obeys the condition to be a representation is close to a unitary representation.

Theorem 14.2. *Let G be a finite group and let $f : G \rightarrow U(n)$ be a function such that $\|f(x)f(y) - f(xy)\|_{HS} \leq \epsilon\sqrt{n}$ for every $x, y \in G$. Then there exist m with $(1 - \epsilon^2)n \leq m \leq (1 - 2\epsilon^2)^{-1}n$, an $n \times m$ partial unitary matrix U , and a unitary representation $\rho : G \rightarrow U(m)$, such that*

$$\|f(x) - U\rho(x)U^*\| \leq 31\epsilon\sqrt{n}$$

for every $x \in G$.

15. CONCLUSION

Now that we have seen several different generalizations of Fourier analysis (though not a complete list), we can draw up a checklist of the properties that a generalization is likely to need in order to be useful. Ideally we would have all of the following.

- A Parseval identity
- A convolution identity
- An inversion formula
- A quasirandomness-measuring norm
- An inverse theorem for the quasirandomness-measuring norm

Sometimes we can indeed get all of these, but in situations where we can't, it turns out that just having the last two properties is sufficient for some very interesting applications. And in some cases it remains a fascinating challenge to find new versions of the generalizations with some of the properties that the current versions lack.